

BOUNDING THE AREA OF A CENTERED DUAL TWO-CELL BELOW, GIVEN LOWER BOUNDS ON ITS SIDE LENGTHS

JASON DEBLOIS

ABSTRACT. For a locally finite set \mathcal{S} in the hyperbolic plane, suppose C is a compact, n -edged two-cell of the centered dual complex of \mathcal{S} , a coarsening of the Delaunay tessellation introduced in the author's prior work. We describe an effectively computable lower bound for the area of C , given an n -tuple of positive real numbers bounding its side lengths below, and for $n \leq 9$ implement an algorithm to compute this bound. For geometrically reasonable side-length bounds, we expect the area bound to be sharp or near-sharp.

This paper upgrades the centered dual machine, which the author used in [4] to give sharp upper bounds on the maximal injectivity radius of complete, orientable, finite-area hyperbolic surfaces. Two-cells of the centered dual complex are obtained by grouping Delaunay cells that are not “centered”, with the goal of producing area bounds from side length bounds. (The construction is more thoroughly reviewed in Section 1 below.) Theorem 3.31 of [4], one of the main results of that paper, realizes this goal. It gives a lower bound in terms of $d > 0$, on the area of an arbitrary centered dual two-cell with all edge lengths at least d .

The main result of this paper generalizes and strengthens that one. Below for a tree T we refer by the *frontier* of T to the collection of edges of some ambient graph that intersect T but do not lie in it; we assume that each frontier edge has exactly one vertex in T .

Theorem 2.11. *Let C be a compact two-cell of the centered dual complex of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$ such that for some $\mathbf{b} = (b_1, \dots, b_n) \in (\mathbb{R}^+)^n$ and enumeration of the edges of C , the i^{th} edge has length at least b_i for each i . Then $\text{area}(C) \geq \min\{B_T(\sigma(\mathbf{b})) \mid T \in \mathcal{T}_n, \sigma \in S_n\}$, where B_T is the area bounding function defined in Proposition 2.8, S_n is the symmetric group on n letters, $\sigma \in S_n$ acts on \mathbf{b} by permutation of entries, and \mathcal{T}_n is the collection of compact, rooted trees T with frontier \mathcal{F} of order n and each vertex trivalent in $V = T \cup \bigcup_{f \in \mathcal{F}} f$.*

For a self-contained but necessarily more elaborate statement of this result see Corollary 3.4. Theorem 2.11 generalizes [4, Thrm. 3.31] by allowing different bounds for the lengths of different edges of C . Even when all edge length bounds are the same, the area bound offered by Theorem 2.11 is stronger than that of its predecessor for $n > 4$ (it is identical for $n = 3$ or 4). See Proposition 3.8. In fact, we expect it to be sharp for “geometrically reasonable” edge length bounds, see Remark 3.11.

I intend to use this in the future to study arc length spectra of hyperbolic surfaces. I used Theorem 3.31 of [4] there to prove that paper's sharp upper bound on maximal injectivity radius; or, equivalently, on the length of the shortest non-constant geodesic arc based at a point p (by which I mean one with both endpoints at p , but possibly a corner there).

Theorem 2.11 can be used analogously to bound, say, the length of the second-shortest such arc as a function of the length of the shortest. Problems of this nature arise naturally when studying hyperbolic three-manifolds with totally geodesic boundary; see eg. [8].

The additional strength and generality of Theorem 2.11 comes at considerable computational expense. Whereas the bound of [4, Thrm. 3.31] is given by a formula requiring essentially a single computation, evaluating the bound here for a two-cell with n edges requires performing three computational tasks: enumerating S_n , enumerating \mathcal{T}_n , and evaluating $B_T(\sigma(\mathbf{b}))$ for each one. The first problem alone has complexity which is at least factorial in n .

Section 3 describes a Python module, *minimizer.py* containing a script *minimize()* that computes the bounds of Theorem 2.11. To do so it calls an existing Python script, *iter-tools.permutations()*, for enumerating permutations; a script *treecrawler()* (also in *minimizer.py*) that computes $B_T(\mathbf{b})$, given T and \mathbf{b} ; and a hand-compiled library, *forest.txt*, of trees in \mathcal{T}_n for $n \leq 9$. So it can actually only compute bounds for cells with up to nine edges.

Remark 0.1. The ancillary materials include *minimizer.py* and *forest.txt*. After downloading them only one modification is required to run *minimizer.minimize()* in a Python 2.7. n interpreter (and possibly others). See the beginning of Section 3.

It is certainly possible to write an algorithm to enumerate \mathcal{T}_n for arbitrary n , and hence to remove the limitation to $n \leq 9$. However enumerating it without redundancy seems more involved, so in any case it is useful to have a classification in low complexity (see Figure 3.2). And this is enough for our purpose here, which is just to get a sense for how the bounds of Theorem 2.11 behave. We carry this out by exploring a few examples in Section 3.3.

Section 1 introduces the centered dual decomposition and establishes notation. We prove Theorem 2.11 in Section 2 by deepening some aspects of the argument in [4]. Of particular note, Corollary 2.7 significantly improves Proposition 3.23 of [4], a key result limiting which points in the “admissible space” $\overline{Ad}(\mathbf{d}_{\mathcal{F}})$ can minimize the area function D_T .

1. THE GEOMETRIC AND CENTERED DUAL DECOMPOSITIONS

Here we will give a brief, conceptual introduction to the subject of this paper, culminating in a description of compact two-cells of the centered dual decomposition determined by a finite subset of a hyperbolic surface. The picture we describe here is fully fleshed out in [4], and we refer the reader there for details, proofs, and the general case. In subsection 1.1 we establish notation that we will use in the remainder of the paper.

Suppose \mathcal{S} is a locally finite subset of \mathbb{H}^2 . The *Voronoi tessellation* of \mathcal{S} is a locally finite convex polygonal decomposition of \mathbb{H}^2 with two-cells in bijective correspondence with \mathcal{S} . For each $s \in \mathcal{S}$, the Voronoi two-cell containing s is

$$V_s = \{x \in \mathbb{H}^2 \mid d(x, s) \leq d(x, s') \text{ for all } s' \in \mathcal{S}\}.$$

Each Voronoi vertex v is of the form $\bigcap_{i=1}^n V_{s_i}$ for a finite collection $\{s_i\} \subset \mathcal{S}$ such that $d(s_i, v) \equiv J$ is minimal among all $s \in \mathcal{S}$. The *geometric dual two-cell* dual to v is the convex




FIGURE 1.1. The possible five-edged centered dual two-cells, outlined in red, and for each its dual tree T (black, bold) and the frontier of T (black, dashed).

hull of the s_i . It is *cyclic*; ie. inscribed in a circle, its *circumcircle*, which has radius J and center v . See [3, §5], and Theorem 5.9 there in particular, and cf. [4, §1].

We say a geometric dual two-cell is *centered* if the center of its circumcircle (ie. its dual Voronoi vertex) is contained in its interior. A Voronoi edge e is centered if it intersects its geometric dual edge (which joins s to s' if $e = V_s \cap V_{s'}$) in its interior; if e is not centered we orient it pointing away from its geometric dual edge. The two notions of centeredness are related: if C is a non-centered geometric dual two-cell then its dual Voronoi vertex is the initial point of a non-centered Voronoi edge, and the geometric dual to the initial vertex of every non-centered Voronoi edge is non-centered [4, Lemma 2.5].

We use components of the union of non-centered Voronoi edges to organize centered dual two-cells. A compact such component is a finite, rooted tree T with all edges pointing toward its root vertex v_T . Circumcircle radius increases in the orientation direction of non-centered edges [4, L. 2.3], so v_T is also characterized as the vertex whose geometric dual two-cell has maximal circumcircle radius. See Lemma 2.7, Definition 2.8, and Proposition 2.9 of [4].

A compact centered dual two-cell C is either the geometric dual to a single Voronoi vertex contained in only centered edges (the *centered case*), or it is the union of geometric duals to vertices of a compact component of the union of non-centered Voronoi edges. In this case we say C is dual to T (see [4, Definition 2.11]); in the centered case we say C is dual to $T = \{v\}$, its dual Voronoi vertex.

The *frontier* of a component T of the union of non-centered Voronoi edges is the set of (e, v) such that e is a Voronoi edge not contained in T and $v \in e \cap T$ is a vertex of e . (So if both vertices of $e \not\subset T$ are in T then e contributes twice to the frontier of T .) For such a tree T , the *edge set* of the centered dual two-cell C dual to T is the collection of geometric duals to Voronoi edges contributing to the frontier of T , counted with multiplicity. Figure 1.1 illustrates all combinatorial possibilities for compact centered dual two-cells with five edges.

A convex cyclic polygon is determined up to isometry by its set of edge lengths [5, Prop. 1.8], but this is not true of a centered dual two-cell C that is dual to a component T of the union of non-centered Voronoi edges. However, the geometry of C is *constrained* by its set of edge lengths. The strategy of [4] is to abstract these constraints, defining an *admissible space* parametrizing all possibilities for a two-cell C with fixed combinatorics and edge length collection. The areas of such possibilities are measured by a continuous function on the admissible space, and we produce area bounds by analyzing minima of this function.

1.1. Notation. The notation we use for abstracting the study of geometric dual two-cells was introduced in Section 3.2 of [4]. We let T denote a rooted tree and call its root vertex v_T . We always implicitly regard T as embedded in some ambient graph in which each vertex v of T has valence $n_v \geq 3$, and we take the *frontier* $\mathcal{F} = \{f_1, \dots, f_n\}$ of T to be the collection of edges of this ambient graph that intersect T but do not lie in it.

We also implicitly assume that each frontier edge f_i has exactly one vertex in T . In the geometric context this could fail; that is, there may exist a component T of the union of non-centered Voronoi edges and a Voronoi edge e with distinct vertices v and w , such that $e \cap T = \{v, w\}$. If so then when passing to the abstract context we would denote the edge-vertex pairs (e, v) and (e, w) as f_i and f_j , respectively, for some $i \neq j$.

We denote the edge set of T by \mathcal{E} and study tuples $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$, where $\mathbf{d}_{\mathcal{E}} \in (\mathbb{R}^+)^{\mathcal{E}}$ and $\mathbf{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$ are collections of positive real numbers indexed by \mathcal{E} and \mathcal{F} , respectively. For a vertex v of T contained in edges $e_1, \dots, e_{n_v} \in \mathcal{E} \cup \mathcal{F}$ and any such \mathbf{d} , we let $P_v(\mathbf{d})$ denote the n_v -tuple $(d_{e_1}, \dots, d_{e_{n_v}})$ of entries of \mathbf{d} .

The idea here is that for a centered dual two-cell C dual to T , $\mathbf{d}_{\mathcal{F}}$ records the set of its edge lengths, since boundary edges of C are dual to frontier edges of T by [4, Dfn. 2.11]. (Here if both f_i and f_j correspond to a single Voronoi edge as above we take $d_{f_i} = d_{f_j}$ to be the length of the geometric dual to e .) And $\mathbf{d}_{\mathcal{E}}$ records the set of lengths of geometric dual edges *internal* to C ; ie. edges of intersection between pairs of geometric dual cells contained in C . For such a geometric dual cell with dual Voronoi vertex v , $P_v(\mathbf{d})$ records its edge length collection, where $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$.

Definition 3.10 of [4] describes the admissible space $Ad_T(\mathbf{d}_{\mathcal{F}})$ of a given T and $\mathbf{d}_{\mathcal{F}}$. For a centered dual two-cell C dual to T with edge length collection $\mathbf{d}_{\mathcal{F}}$, with $\mathbf{d}_{\mathcal{E}}$ is produced as above, Lemma 3.14 there shows that $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$ lies in $Ad_T(\mathbf{d}_{\mathcal{F}})$. It is more convenient in practice to deal with a compact space $\overline{Ad}_T(\mathbf{d}_{\mathcal{F}})$ containing $Ad(\mathbf{d}_{\mathcal{F}})$, which is defined in [4, Dfn. 3.15]. We reproduce this below. There for $v \in T^{(0)} - \{v_T\}$, let e_v be the initial edge of the edge arc joining v to v_T .

Definition 1.1 ([4], Definition 3.15). For $\mathbf{d}_{\mathcal{F}} = (d_e \mid e \in \mathcal{F}) \in (\mathbb{R}^+)^{\mathcal{F}}$ let $\overline{Ad}(\mathbf{d}_{\mathcal{F}})$ consist of those $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$, for $\mathbf{d}_{\mathcal{E}} \in (\mathbb{R}^+)^{\mathcal{E}}$, such that:

- (1) For $v \in T^{(0)} - \{v_T\}$, $P_v(\mathbf{d}) \in \mathcal{AC}_{n_v} - \mathcal{C}_{n_v}$ has largest entry d_{e_v} .
- (2) $P_{v_T}(\mathbf{d}) \in \mathcal{C}_{n_T} \cup \mathcal{BC}_{n_T}$, where we refer by n_T to the valence n_{v_T} of v_T in V .
- (3) $J(P_v(\mathbf{d})) \geq J(P_w(\mathbf{d}))$ for each $v \in T^{(0)}$ and $w \in v - 1$, where $v - 1$ is the set of vertices $w \in T^{(0)} - \{v_T, v\}$ such that $v \in e_w$.

Here $\mathcal{AC}_n \subset (\mathbb{R}^+)^n$ is the open set parametrizing cyclic n -gons by their side lengths, see [5, Corollary 1.10]. Its subsets \mathcal{C}_n and \mathcal{BC}_n respectively parametrize *centered* and *semicyclic* n -gons, those with circumcircle centers in their interiors or, respectively, in a side. See Propositions 1.11 and 2.2 of [5]. Condition (3) above refers to the function $J: \mathcal{AC}_n \rightarrow \mathbb{R}^+$ that records circumcircle radius of cyclic n -gons [5, Prop. 1.14].

Remark 1.2. Below at times we will take $T = \{v_T\}$ (the centered case) as the base case of an inductive argument. In this case since \mathcal{E} is empty we omit $\mathbf{d}_{\mathcal{E}}$. Conditions (1) and (3) of Definition 1.1 hold vacuously, so appealing to condition (2) we take $\overline{Ad}(\mathbf{d}_{\mathcal{F}}) = \{\mathbf{d}_{\mathcal{F}}\}$ if $P_{v_T}(\mathbf{d}_{\mathcal{F}}) \in \mathcal{C}_{n_T} \cup \mathcal{BC}_{n_T}$ and $\overline{Ad}(\mathbf{d}_{\mathcal{F}}) = \emptyset$ otherwise.

Definition 3.13 of [4] introduces the *area function* $D_T(\mathbf{d}) = \sum_{v \in T^{(0)}} D_0(P_v(\mathbf{d}))$, where D_0 from [5, Cor. 1.17] is the function $\mathcal{AC}_n \rightarrow \mathbb{R}$ that measures the area of cyclic n -gons [5, L. 2.1]. D_T measures area of centered dual two-cells [4, L. 3.14]. It is continuous on $\overline{Ad}_T(\mathbf{d}_{\mathcal{F}})$ [4, L. 3.22], and our primary aim here is to understand its minimizers and minima on this set.

2. THEORY

In this section we will prove Theorem 2.11, the generalization of Theorem 3.31 of [4] described in the introduction. We follow the broad strokes of the approach in [4]. First, in Section 2.1 we prove Corollary 2.7, an upgrade of [4, Prop. 3.23] that characterizes minimizers of D_T on $\overline{Ad}_T(\mathbf{d}_{\mathcal{F}})$ for a given $\mathbf{d}_{\mathcal{F}}$. We apply this in Section 2.2 to prove Proposition 2.8, an improvement on [4, Prop. 3.30] that bounds $D_T(\mathbf{d})$ below by a function $B_T(\mathbf{b}_{\mathcal{F}})$ for all $\mathbf{d} \in \overline{Ad}_T(\mathbf{d}_{\mathcal{F}})$ where each entry of $\mathbf{d}_{\mathcal{F}}$ is bounded below by the corresponding entry of $\mathbf{b}_{\mathcal{F}}$.

The results above apply to a fixed rooted tree T . We prove Theorem 2.11 by simply minimizing over all trees and using one new ingredient: Lemma 2.10, which allows us to reduce to the trivalent case. This Lemma reverses the action of Lemma 3.28 of [4], showing for some trees T that $\overline{Ad}_T(\mathbf{d}_{\mathcal{F}}) \subset \overline{Ad}_{T_0}(\mathbf{d}_{\mathcal{F}_0})$ for a related tree T_0 with *more* edges than T .

2.1. Minimizers on $\overline{Ad}_T(\mathbf{d}_{\mathcal{F}})$. Proposition 3.23 of [4] supplies a key tool for giving lower bounds on areas of centered dual two-cells, asserting that each minimum (in fact each local minimum) point of D_T on $\overline{Ad}(\mathbf{d}_{\mathcal{F}})$ satisfies one of three criteria listed there. Here we will more closely analyze the situations described there.

Lemma 2.1. *For a compact rooted tree $T \subset V$ with root vertex v_T , edge set \mathcal{E} , and frontier $\mathcal{F} = \{f_0, \dots, f_{k-1}\}$, given $\mathbf{d}_{\mathcal{F}} = (d_{f_0}, \dots, d_{f_{k-1}}) \in (\mathbb{R}^+)^{\mathcal{F}}$ and $\mathbf{d}_{\mathcal{E}} \in (\mathbb{R}^+)^{\mathcal{E}}$ let $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$. Suppose such a tuple \mathbf{d} has the following properties:*

- (1) *For each $v \in T^{(0)} - \{v_T\}$, $P_v(\mathbf{d}) \in \mathcal{AC}_{n_v} - (\mathcal{C}_{n_v} \cup \mathcal{BC}_{n_v})$ has largest entry d_{e_v} , where n_v is the valence of v in V and e_v is the initial edge of the arc in T joining v to v_T ;*
- (2) *$P_{v_T}(\mathbf{d}) \in \mathcal{AC}_{n_T}$, where $n_T = n_{v_T}$, and if $P_{v_T} \notin \mathcal{C}_{n_T}$ then $e \in \mathcal{F}$ if d_e is maximal among entries of $P_{v_T}(\mathbf{d})$; and*
- (3) *$J(P_v(\mathbf{d})) = J(P_{v_T}(\mathbf{d}))$ for each $v \in T^{(0)}$, where $J(P_v(\mathbf{d}))$ is the circumcircle radius of the cyclic polygon with edge length collection $P_v(\mathbf{d})$.*

Then $\mathbf{d}_{\mathcal{F}}$ is in \mathcal{AC}_k , and in \mathcal{C}_k or \mathcal{BC}_k if and only if $P_{v_T}(\mathbf{d})$ is in \mathcal{C}_{n_T} or \mathcal{BC}_{n_T} , respectively. If $\mathbf{d}_{\mathcal{F}} \notin \mathcal{C}_k$ then there is a unique $f_i \in \mathcal{F}$ such that d_{f_i} is maximal, with $v_T \in f_i$ and $d_{f_i} > d_e$ for all $e \in (\mathcal{E} \cup \mathcal{F}) - \{f_i\}$. Also, $J(\mathbf{d}_{\mathcal{F}}) = J(P_v(\mathbf{d}))$ for each $v \in T^{(0)}$, and

$$(2.1.1) \quad D_0(\mathbf{d}_{\mathcal{F}}) = \sum_{v \in T^{(0)}} D_0(P_v(\mathbf{d})).$$

Moreover, a cyclic k -gon with side length collection $\mathbf{d}_{\mathcal{F}}$ is tiled by cyclic n_v -gons P_v with side length collections $P_v(\mathbf{d})$, where v runs over all vertices of T .

Proof. Note that if $T = \{v_T\}$ then hypotheses (1) and (3) hold vacuously, and the result is a tautology. Below we will first address the case that T has one edge, then prove the general case by induction. The one-edge case is an altered version of Lemma 3.25 of [4], with a stronger conclusion and a subtly stronger hypothesis. That result takes as input an m -tuple $\mathbf{c}_0 = (c_0, \dots, c_{m-1})$ and an n -tuple $\mathbf{d}_0 = (d_0, \dots, d_{n-1})$. We will apply it here in the one-edged case with $\mathbf{d}_0 = P_{v_T}(\mathbf{d})$ and $\mathbf{c}_0 = P_v(\mathbf{d})$, where v is the other vertex of T . Our hypothesis (3) above implies the first bulleted hypothesis there, that $J(\mathbf{c}_0) = J(\mathbf{d}_0)$; (1) implies the second bullet there with d_{e_v} here in the role of $d_0 = c_0$ there, where e_v is the sole edge of T ; and (2) here implies the third. In the notation of [4, L. 3.25], (2) in fact asserts:

- $\mathbf{d}_0 \in \mathcal{AC}_n$, and if $\mathbf{d}_0 \notin \mathcal{C}_n$ then d_0 is not maximal among the d_i .

This excludes the possibility that $\mathbf{d}_0 \in \mathcal{BC}_n$ has maximal entry d_0 , which was allowed in the third bulleted hypothesis of [4, L. 3.25]. In any case, since its hypotheses are satisfied the proof and conclusions of that result hold.

It asserts that $\mathbf{d} \doteq (c_1, \dots, c_{m-1}, d_1, \dots, d_{n-1})$ is in \mathcal{AC}_{m+n-2} . Note that since $\mathcal{E} = \{e_v\}$ has only one element, each of $\mathbf{d}_0 = P_{v_T}(\mathbf{d})$ and $\mathbf{c}_0 = P_v(\mathbf{d})$ has all of its entries but $d_{e_v} = d_0 = c_0$ from $\mathbf{d}_{\mathcal{F}}$, so $\mathbf{d}_{\mathcal{F}} = \mathbf{d}$ and $k = m + n - 2$. The conclusion of Lemma 3.25 of [4] thus asserts in our terms that $\mathbf{d}_{\mathcal{F}}$ is in \mathcal{AC}_k , and in $\mathcal{C}_k \cup \mathcal{BC}_k$ if and only if $P_{v_T}(\mathbf{d}) \in \mathcal{C}_{n_T} \cup \mathcal{BC}_{n_T}$; and that $D_0(\mathbf{d}_{\mathcal{F}}) = D_0(P_v(\mathbf{d})) + D_0(P_{v_T}(\mathbf{d}))$.

For the one-edged case of our result we must strengthen this conclusion with four additional assertions (assuming the bulleted hypothesis above). We claim first that $\mathbf{d}_{\mathcal{F}} = \mathbf{d}$ is the side length collection of a cyclic polygon P tiled by polygons P_v and P_{v_T} with respective side length collections $P_v(\mathbf{d}) = \mathbf{c}_0$ and $P_{v_T}(\mathbf{d}) = \mathbf{d}_0$. This is in fact recorded in the proof of [4, Lemma 3.25], where P_v is called P_0 and P_{v_T} is Q_0 . It implies our second additional assertion, that $J(\mathbf{d}_{\mathcal{F}}) = J(P_v(\mathbf{d})) = J(P_{v_T}(\mathbf{d}))$, since P shares a circumcircle with P_v and P_{v_T} .

We also need that if $\mathbf{d}_{\mathcal{F}} \notin \mathcal{C}_k$ then its unique largest entry comes from among the d_{f_i} with $v_T \in f_i$, i.e. from among the $d_i \neq d_0$, in the language of [4, L. 3.25]. The proof there shows that $P_0 \cap Q_0 = \gamma_0$ is a side of each with length $c_0 = d_0 = d_{e_v}$, and that P_0 and the circumcircle center v lie in opposite half-spaces bounded by the geodesic containing γ_0 . Proposition 2.2 of [5] then implies that γ_0 is unique with this property among sides of P_0 . If $\mathbf{d}_{\mathcal{F}} \notin \mathcal{C}_k$, i.e. $\mathbf{d} \notin \mathcal{C}_{m+n-2}$, then the unique longest side γ_{i_0} of P is characterized by the fact that v and P lie in opposite half-spaces bounded by the geodesic containing γ_{i_0} , again by [5, Prop. 2.2]. This implies γ_{i_0} is a side of Q_0 other than γ_0 , since these comprise the remaining sides of P , and this assertion follows.

We finally require that $\mathbf{d}_{\mathcal{F}}$ is in \mathcal{C}_k or \mathcal{BC}_k if and only if $P_{v_T}(\mathbf{d}) = \mathbf{d}_0$ is in \mathcal{C}_n or \mathcal{BC}_n respectively. It is asserted in [4, L. 3.25] that $\mathbf{d} \in \mathcal{C}_{m+n-2} \cup \mathcal{BC}_{m+n-2}$ if and only if $\mathbf{d}_0 \in \mathcal{C}_n \cup \mathcal{BC}_n$. We have the additional fact, again by [5, Prop. 2.2], that $\mathbf{d} \in \mathcal{BC}_{m+n-2}$ if and only if v lies in a side of P . If this is so then by the above v lies in a side of Q_0 , so $\mathbf{d}_0 \in \mathcal{BC}_n$.

On the other hand, our strengthened hypothesis does not allow $v \in \gamma_0$, so if $\mathbf{d}_0 \in \mathcal{BC}_n$ then v lies in a side of Q_0 that is a side of P . Hence $\mathbf{d} \in \mathcal{BC}_{m+n-2}$, and the second assertion is proved. This gives the one-edge case of the current result.

We now proceed to the inductive step. Let T be a compact rooted tree with at least two edges, and let v_0 be a vertex farthest from v_T in T . Then v_0 is contained in a single edge e_{v_0} of T , and we take $T_0 = \overline{T - e_{v_0}}$. Listing the edges containing v_0 as $e_{v_0}, f_{i_1}, \dots, f_{i_{n_{v_0}-1}}$, where all $f_{i_j} \in \mathcal{F}$, the frontier of T_0 in V is $\mathcal{F}_0 = (\mathcal{F} - \{f_{i_j}\}) \cup \{e_{v_0}\}$ and the edge set is $\mathcal{E} - \{e_{v_0}\}$. Given $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$ we obtain $\mathbf{d}_0 = (\mathbf{d}_{\mathcal{E}_0}, \mathbf{d}_{\mathcal{F}_0})$ by omitting the entries d_{f_j} of $\mathbf{d}_{\mathcal{F}}$ and shuffling $d_{e_{v_0}}$ from $\mathbf{d}_{\mathcal{E}}$ to $\mathbf{d}_{\mathcal{F}_0}$. Then $P_v(\mathbf{d}_0) = P_v(\mathbf{d})$ for all $v \in T_0^{(0)}$.

In particular, if T and \mathbf{d} satisfy (1)–(3) then so do T_0 and \mathbf{d}_0 . We suppose this is so, and assume by induction that the desired conclusion holds for T_0 and \mathbf{d}_0 . We now note that the hypotheses of our strengthened [4, Lemma 3.25] (from the one-edged case) are satisfied with \mathbf{c}_0 the n_{v_0} -tuple of dual lengths to the edges containing v_0 and $\mathbf{d}_0 = \mathbf{d}_{\mathcal{F}_0}$, ordered so that $c_0 = d_0 = d_{e_{v_0}}$. Applying that result and noting that \mathbf{d} as described there is $\mathbf{d}_{\mathcal{F}}$, we thus conclude that $\mathbf{d}_{\mathcal{F}} \in \mathcal{AC}_n$; it is in \mathcal{C}_n or \mathcal{BC}_n if and only if $P_{v_T}(\mathbf{d})$ is in \mathcal{C}_{n_T} or \mathcal{BC}_{n_T} , respectively; that $J(\mathbf{d}_{\mathcal{F}}) = J(P_v(\mathbf{d}))$ for all $v \in T^{(0)}$; and that $D_0(\mathbf{d}_{\mathcal{F}}) = \sum_{v \in T^{(0)}} D_0(P_v(\mathbf{d}))$.

Moreover by induction the cyclic n_0 -gon Q_0 of [4, L. 3.25], with side length collection $\mathbf{d}_{\mathcal{F}_0}$, is itself tiled by copies of the P_v for all $v \in T_0^{(0)}$. And since P_0 from [4, L. 3.25] shares the side length collection of P_{v_0} it is isometric to it; hence a cyclic n -gon with side length collection $\mathbf{d}_{\mathcal{F}}$ is tiled by copies of the P_v as desired. \square

Lemma 2.2. *With the hypotheses of Lemma 2.1, for any continuous map $t \mapsto \mathbf{d}_{\mathcal{F}}(t) \in (\mathbb{R}^+)^{\mathcal{F}}$ with $\mathbf{d}_{\mathcal{F}}(0) = \mathbf{d}_{\mathcal{F}}$ there is a continuous map $t \mapsto \mathbf{d}_{\mathcal{E}}(t) \in (\mathbb{R}^+)^{\mathcal{E}}$ on $[0, \epsilon)$ for some $\epsilon > 0$, with $\mathbf{d}_{\mathcal{E}}(0) = \mathbf{d}_{\mathcal{E}}$, such that $\mathbf{d}(t) = (\mathbf{d}_{\mathcal{F}}(t), \mathbf{d}_{\mathcal{E}}(t))$ has properties (1) - (3) of the lemma.*

Proof. We use the assertion in Lemma 2.1 that for $\mathbf{d} = (\mathbf{d}_{\mathcal{F}}, \mathbf{d}_{\mathcal{E}})$ satisfying its hypotheses, a cyclic n -gon P with side length collection $\mathbf{d}_{\mathcal{F}}$ is tiled by copies of P_v for $v \in T^{(0)}$. Thus for each $e \in \mathcal{E}$ and each vertex v of e , the corresponding edge of $P_v(\mathbf{d})$ is a diagonal of P . So its length d_e is given by $\ell_{ij}(\mathbf{d}_{\mathcal{F}})$ for some fixed i and j between 0 and $n - 1$, where ℓ_{ij} is the diagonal-length function of [5, Corollary 1.15].

Now any continuous deformation $\mathbf{d}_{\mathcal{F}}(t)$ of $\mathbf{d}_{\mathcal{F}}$ remains in \mathcal{AC}_n for small t , since \mathcal{AC}_n is open in $(\mathbb{R}^+)^n$. We define $\mathbf{d}_{\mathcal{E}}(t)$ by taking $d_e(t) = \ell_{ij}(\mathbf{d}_{\mathcal{F}}(t))$ for each $e \in \mathcal{E}$, where ℓ_{ij} is the diagonal-length function described above. Then by the definition of the $\ell_{i,j}$, $P_v(\mathbf{d}(t)) \in \mathcal{AC}_{n_v}$ and $J(P_v(\mathbf{d}(t))) = J(P_{v_T}(\mathbf{d}(t))) = J(\mathbf{d}_{\mathcal{F}}(t))$ for all t such that $\mathbf{d}_{\mathcal{F}}(t) \in \mathcal{AC}_n$. Criteria (1) and (2) of the lemma involve only open conditions and so hold for t small enough. \square

Proposition 2.3. *For a compact rooted tree $T \subset V$ with root vertex v_T , edge set \mathcal{E} and frontier $\mathcal{F} = \{f_0, \dots, f_{n-1}\}$, and $\mathbf{d}_{\mathcal{F}} = (d_{f_0}, \dots, d_{f_{n-1}}) \in (\mathbb{R}^+)^{\mathcal{F}}$, suppose a local minimum of $\mathbf{d}_{\mathcal{E}} \mapsto D_T(\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$ on $\overline{Ad}(\mathbf{d}_{\mathcal{F}})$ occurs at $\mathbf{d}_{\mathcal{E}}$ with the following property: for $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$ there exists $v_0 \in T^{(0)}$ such that $J(P_v(\mathbf{d})) = J(P_{v_0}(\mathbf{d}))$ for the terminal vertex v of e_{v_0} . Then for the maximal subtree T_0 of T containing v_0 such that $J(P_v(\mathbf{d})) = J(P_{v_0}(\mathbf{d}))$ for all $v \in T_0^{(0)}$:*

- $v_T \in T_0$;

- $P_{v_T}(\mathbf{d}) \in \mathcal{BC}_{n_T}$, where v_T has valence n_T in V ; and
- for $e_T \in \mathcal{E} \cup \mathcal{F}$ containing v_T such that d_{e_T} is maximal among all such edges, either $e_T \in \mathcal{F}$ or e_T is an edge of T_0 and its other endpoint v'_T is on the boundary of T_0 ; ie. it is of valence one in T_0 .

Moreover, if $e_T \in \mathcal{E}$ then for each $v \in v'_T - 1$, $P_v(\mathbf{d}) \in \mathcal{BC}_{n_v}$.

Recall that Definition 1.1(1) asserts for each $v \in T^{(0)} - \{v_T\}$ that $P_v(\mathbf{d})$ is in $\mathcal{AC}_{n_v} - \mathcal{C}_{n_v}$. We begin by noting a stronger fact for vertices of T_0 .

Lemma 2.4. *With hypotheses and notation as in Proposition 2.3, let v_1 be the nearest vertex of T_0 to v_T in T (so $v_T \in T_0 \Leftrightarrow v_T = v_1$). If $v_1 \neq v_T$ then for each $v \in T_0^{(0)} - \{v_1\}$, $P_v(\mathbf{d}) \in \mathcal{AC}_{n_v} - (\mathcal{C}_{n_v} \cup \mathcal{BC}_{n_v})$. If $v_1 = v_T$ then this still holds for all but at most one $v \in T_0^{(0)} - \{v_1\}$. If in this case there does exist $v \in T_0^{(0)} - \{v_1\}$ with $P_v(\mathbf{d}) \in \mathcal{BC}_{n_v}$, then e_v joins v to v_T , $P_{v_T}(\mathbf{d}) \in \mathcal{BC}_{n_T}$, and both $P_{v_T}(\mathbf{d})$ and $P_v(\mathbf{d})$ have maximal entry d_{e_v} .*

Proof. For v as above we just need to show that $P_v(\mathbf{d}) \notin \mathcal{BC}_{n_v}$. Note that for such v , $P_v(\mathbf{d})$ has d_{e_v} as its largest entry, by Definition 1.1(1), and by Proposition 2.2 of [5], d_{e_v} is *unique* with this property. Since $v \neq v_1$, the other endpoint v' of e_v lies in T_0 , so $J(P_v(\mathbf{d})) = J(P_{v'}(\mathbf{d}))$. If $P_v(\mathbf{d})$ lies in \mathcal{BC}_{n_v} then by [5, Prop. 1.11], $J(P_v(\mathbf{d})) = d_{e_v}/2$. Hence also $J(P_{v'}(\mathbf{d})) = d_{e_v}/2$, so by the same result $P_{v'}(\mathbf{d}) \in \mathcal{BC}_{n_{v'}}$ has largest entry d_{e_v} .

This is a contradiction if $v' \neq v_T$, since then it has unique largest entry $d_{e_{v'}} > d_{e_v}$ by Definition 1.1(1). If $v' = v_T$ then since $v' \in T_0$ we must have $v_1 = v_T$, and the Lemma follows from the previous paragraph. \square

For the sake of readability we will prove the first two assertions of the Proposition separately.

Lemma 2.5. *With the hypotheses and notation of Proposition 2.3, $v_T \in T_0$ and $P_{v_T} \in \mathcal{BC}_{n_T}$.*

Proof. Suppose that either v_T is not in T_0 (ie. $v_1 \neq v_T$), or $v_T \in T_0$ ($v_1 = v_T$) and $P_{v_T}(\mathbf{d}) \in \mathcal{C}_{n_T}$. So in particular, $P_v(\mathbf{d}) \notin \mathcal{BC}_{n_v}$ for all $v \in T_0^{(0)} - \{v_1\}$ by Lemma 2.4. Without loss of generality assume v_0 is a farthest vertex of T_0 from v_T , and refer by e_0 to e_{v_0} . Let $T_1 = T_0 - (\text{int}(e_0) \cup v_0)$, and take v_1 as its root vertex. Let \mathcal{F}_1 be the frontier of T_1 in V and name its edge set \mathcal{E}_1 , and let $\mathbf{d}_1 = (\mathbf{d}_{\mathcal{E}_1}, \mathbf{d}_{\mathcal{F}_1})$ take its entries from \mathbf{d} . Define a deformation $\mathbf{d}_{\mathcal{F}_1}(t)$ of $\mathbf{d}_{\mathcal{F}_1}$ by taking $d_e(t) \equiv d_e$ for all $e \in \mathcal{F}_1 - \{e_0\}$ and $d_{e_0}(t) = d_{e_0} - t$. We will show that this determines a deformation $\mathbf{d}(t) \in \overline{\text{Ad}}(\mathbf{d}_{\mathcal{F}})$ such that $D_T(\mathbf{d}(t))$ is decreasing.

We first note that \mathbf{d}_1 satisfies criteria (1)–(3) of Lemma 2.1: property (1) is Lemma 2.4, and (3) is inherited from T_0 . If $v_1 \neq v_T$ then criterion (2) follows from the facts that $P_{v_1}(\mathbf{d}_1) = P_{v_1}(\mathbf{d}) \in \mathcal{AC}_{n_{v_1}} - \mathcal{C}_{n_{v_1}}$ has maximal entry $d_{e_{v_1}}$, and $e_{v_1} \in \mathcal{F}_1$ since by construction v_1 is nearest v_T in T_1 . If $v_1 = v_T$ then by hypothesis $P_{v_T}(\mathbf{d}) \in \mathcal{C}_{n_T}$, and criterion (2) is immediate.

Lemma 2.1 now implies that if $v_1 \neq v_T$ and hence $P_{v_1}(\mathbf{d}_1) \in \mathcal{AC}_{n_{v_1}} - \mathcal{C}_{n_{v_1}}$, then $\mathbf{d}_{\mathcal{F}_1} \in \mathcal{AC}_{n_1} - \mathcal{C}_{n_1}$, where $n_1 = |\mathcal{F}_1|$, and it has maximal entry $d_{e_{v_1}}$. Otherwise $\mathbf{d}_{\mathcal{F}_1} \in \mathcal{C}_{n_1}$. We claim:

2.5.1. If $\mathbf{d}_{\mathcal{F}_1} \in \mathcal{AC}_{n_1} - \mathcal{C}_{n_1}$ then $\mathbf{d}_{\mathcal{F}_1}(t)$ remains in $\mathcal{AC}_{n_1} - \mathcal{C}_{n_1}$ for small $t > 0$; otherwise $\mathbf{d}_{\mathcal{F}_1}(t) \in \mathcal{C}_{n_1}$ for small $t > 0$.

Proof of 2.5.1. This holds if $\mathbf{d}_{\mathcal{F}_1} \in \mathcal{AC}_{n_1} - (\mathcal{C}_{n_1} \cup \mathcal{BC}_{n_1})$, or if $\mathbf{d}_{\mathcal{F}_1} \in \mathcal{C}_{n_1}$, simply because these sets are open. If $\mathbf{d}_{\mathcal{F}_1}$ lies in \mathcal{BC}_{n_1} then Proposition 1.12 of [5] implies that $d_{e_{v_1}} = b_0(d_{e_0}, \dots, d_{e_{n_1-1}})$ for the function b_0 defined there, where we have enumerated \mathcal{F}_1 as $\{e_0, \dots, e_{n_1-1}, e_{v_1}\}$. Since b_0 strictly increases in each variable (also by that result), and $d_{e_0}(t) < d_{e_0}$ for each $t > 0$, it follows that $b_0(d_{e_0}(t), \dots, d_{e_{n_1-1}}(t)) < b_0(d_{e_1}, \dots, d_{e_{n_1-1}})$ for each such t . Then $d_{e_{v_1}}(t) \equiv d_{e_{v_1}}$ exceeds $b_0(d_{e_0}(t), \dots, d_{e_{n_1-1}}(t))$ for each $t > 0$, so the claim follows in this case from [5, Cor. 4.10]. \square

By Lemma 2.2, $\mathbf{d}_{\mathcal{F}_1}(t)$ determines a deformation $\mathbf{d}_{\mathcal{E}_1}(t)$ of $\mathbf{d}_{\mathcal{E}_1}$ for small $t \geq 0$ such that properties (1) - (3) of Lemma 2.1 continue to hold for $\mathbf{d}_1(t) = (\mathbf{d}_{\mathcal{E}_1}(t), \mathbf{d}_{\mathcal{F}_1}(t))$. We extend $\mathbf{d}_1(t)$ to $\mathbf{d}(t) = (\mathbf{d}_{\mathcal{E}}(t), \mathbf{d}_{\mathcal{F}}(t))$ by taking $d_e(t) \equiv d_e$ for each $e \in \mathcal{E} \cup \mathcal{F} - (\mathcal{E}_1 \cup \mathcal{F}_1)$. In particular $\mathbf{d}_{\mathcal{F}}(t) \equiv \mathbf{d}_{\mathcal{F}}$ since $e_0 \in \mathcal{E}$. We claim that $\mathbf{d}(t) \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$ for small enough $t > 0$.

The only edges of T that change length under $\mathbf{d}(t)$ are edges of T_0 , so if $P_v(\mathbf{d})$ changes under $\mathbf{d}(t)$ then $v \in T_0$. Property (1) of Definition 1.1 is thus immediate for $v \in T^{(0)} - T_0^{(0)}$. It follows for $v \in T_0^{(0)} - \{v_1\}$ from Lemma 2.4 and the fact that $\mathcal{AC}_{n_v} - (\mathcal{C}_{n_v} \cup \mathcal{BC}_{n_v})$ is open in $(\mathbb{R}^+)^{n_v}$, and, in the case $v_1 \neq v_T$, for v_1 by combining 2.5.1 above with the first assertion of Lemma 2.1. Property (2) of Definition 1.1, that $P_{v_T}(\mathbf{d}(t)) \in \mathcal{C}_{n_T} \cup \mathcal{BC}_{n_T}$, is immediate if $v_1 \neq v_T$ and otherwise follows from 2.5.1 and Lemma 2.2.

For property (3) of Definition 1.1 we must separately consider several possibilities for $v \in T^{(0)}$ and $w \in v - 1$. If neither v nor w lies in T_0 then we have

$$J(P_v(\mathbf{d}(t))) \equiv J(P_v(\mathbf{d})) \geq J(P_w(\mathbf{d})) \equiv J(P_w(\mathbf{d}(t)))$$

for all t . If v is not in T_0 but w is then $w = v_1 \neq v_T$, and by definition of T_0 the initial inequality is strict: $J(P_v(\mathbf{d})) > J(P_{v_1}(\mathbf{d}))$. So it is preserved for small $t > 0$. The same idea holds if $v \in T_0^{(0)}$ but $w \notin T_0$: the strict initial inequality $J(P_v(\mathbf{d})) > J(P_w(\mathbf{d}))$ is preserved for small $t > 0$. If v and w lie in T_1 —ie. $v, w \in T_0$ and $w \neq v_0$ —then by Lemmas 2.2 and 2.1 $J(P_v(\mathbf{d}(t))) \equiv J(\mathbf{d}_{\mathcal{F}_1}(t)) \equiv J(P_w(\mathbf{d}(t)))$. So the claim is finally proved by establishing property (3) in the case $w = v_0$, so $v \in T_1$. This follows from:

2.5.2. For $t > 0$, $J(\mathbf{d}_{\mathcal{F}_1}(t)) > J(P_{v_0}(\mathbf{d}(t)))$.

Proof of 2.5.2. We have $\frac{d}{dt}J(\mathbf{d}_{\mathcal{F}_1}(t)) = -\frac{\partial}{\partial d_{e_0}}J(\mathbf{d}_{\mathcal{F}_1})$ and $\frac{d}{dt}J(P_{v_0}(\mathbf{d}(t))) = -\frac{\partial}{\partial d_{e_0}}J(P_{v_0}(\mathbf{d}))$ at $t = 0$. Since d_{e_0} is the largest entry of $P_{v_0}(\mathbf{d}) \in \mathcal{AC}_{n_{v_0}} - (\mathcal{C}_{n_{v_0}} \cup \mathcal{BC}_{n_{v_0}})$, Proposition 1.14 of [5] implies that the latter derivative is less than $-1/2$. In the case that $v_1 \neq v_T$, that result implies that $\frac{d}{dt}J(\mathbf{d}_{\mathcal{F}_1}(t)) > 0$. If $v_1 = v_T$ then since $\mathbf{d}_{\mathcal{F}_1} \in \mathcal{C}_{n_1}$ by hypothesis it gives that this derivative is greater than $-1/2$, so in both cases we have the desired inequality. \square

We now show that $D_T(\mathbf{d}(t))$ is decreasing for small $t > 0$. To do so we use Lemma 2.1 (which applies by construction of $\mathbf{d}(t)$) to rewrite $\sum_{v \in T_1^{(0)}} D_0(P_v(\mathbf{d}(t)))$ as $D_0(\mathbf{d}_{\mathcal{F}_1}(t))$, yielding:

$$(2.5.1) \quad D_T(\mathbf{d}(t)) = D_0(\mathbf{d}_{\mathcal{F}_1}(t)) + D_0(P_{v_0}(\mathbf{d}(t))) + \sum_{v \in T^{(0)} - T_0^{(0)}} D_0(P_v(\mathbf{d}(t)))$$

For any $v \in T^{(0)}$, if the edges of $\mathcal{E} \cup \mathcal{F}$ containing v are e_{i_1}, \dots, e_{i_k} (for $k = n_v$) then $P_v(\mathbf{d}(t)) = (d_{e_{i_1}}(t), \dots, d_{e_{i_k}}(t))$ and so by the chain rule we have:

$$(2.5.2) \quad \frac{d}{dt} D_0(P_v(\mathbf{d}(t))) = \sum_{j=1}^k \frac{\partial}{\partial d_{e_{i_j}}} D_0(P_v(\mathbf{d}(t))) \frac{d}{dt} d_{e_{i_j}}(t)$$

If $v \in T^{(0)} - T_0^{(0)}$ this implies in particular that $D_0(P_v(\mathbf{d}(t)))$ is constant, since $d_e(t)$ is constant for each edge e containing such a vertex v . So the rightmost sum of (2.5.1) is constant in t . We now compute $\frac{d}{dt} (D_0(\mathbf{d}_{\mathcal{F}_1}(t)) + D_0(P_{v_0}(\mathbf{d}(t))))$ by applying the chain rule as in (2.5.2) and Proposition 2.3 of [5]. This gives:

$$\sqrt{\frac{1}{\cosh^2(d_{e_0}/2)} - \frac{1}{\cosh^2 J(P_{v_0}(\mathbf{d}(t)))}} - \sqrt{\frac{1}{\cosh^2(d_{e_0}/2)} - \frac{1}{\cosh^2 J(\mathbf{d}_{\mathcal{F}_1}(t))}}$$

By 2.5.2, this quantity is negative for $t > 0$, so indeed $D_T(\mathbf{d}(t))$ is decreasing. The Lemma follows, since by hypothesis \mathbf{d} is a local minimizer for D_T on $\overline{Ad}(\mathbf{d}_{\mathcal{F}})$, and we produced $\mathbf{d}(t)$ assuming that either $v_T \notin T_0$ or $v_T \in T_0$ but $P_{v_t} \in \mathcal{C}_{n_T}$. \square

Proof of Proposition 2.3. For T_0 as described in the Proposition we have $v_T \in T_0$ and $P_{v_T}(\mathbf{d}) \in \mathcal{BC}_{n_T}$, by Lemma 2.5. Let $e_T \in \mathcal{E} \cup \mathcal{F}$ be the edge containing v_T that has d_{e_T} maximal among all such edges. We now suppose by way of contradiction that $e_T \in \mathcal{E}$, and that its other endpoint v'_T is not on the boundary of T_0 . We will show that then \mathbf{d} is still not a local minimum of D_T on $\overline{Ad}(\mathbf{d}_{\mathcal{F}})$.

We begin by observing that Proposition 1.11 of [5] shows that $v'_T \in T_0$ and $P_{v'_T}(\mathbf{d}) \in \mathcal{BC}_{n'_T}$, since $J(P_{v_T}(\mathbf{d})) = d_{e_T}/2$. This is because on the one hand, $J(P_{v'_T}(\mathbf{d})) \leq J(P_{v_T}(\mathbf{d})) = d_{e_T}/2$ by Definition 1.1(3), but on the other $J(P_{v'_T}(\mathbf{d})) \geq d_{e_T}/2$ since the circumradius of a cyclic polygon is at least half of each of its side lengths (cf. [5, Prop. 1.5]).

Let T' be the maximal subtree of T_0 containing v'_T but not v_T , let v'_0 be a farthest vertex of T' from v'_T , and refer by e'_0 to the initial edge $e_{v'_0}$ of the arc joining v'_0 to v'_T . Let $T'_1 = T' - (\text{int}(e'_0) \cup v'_0)$, and let \mathcal{F}'_1 be its frontier in V . We enumerate \mathcal{F}'_1 as $\{e'_0, \dots, e_{n'_1-2}, e_T\}$, so in particular $n'_1 = |\mathcal{F}'_1|$, and define a deformation $\mathbf{d}_{\mathcal{F}'_1}(t)$ of the tuple $\mathbf{d}_{\mathcal{F}'_1}$ that takes its entries from \mathbf{d} as follows: take $d_{e'_0}(t) = d_{e'_0} - t$, let $d_{e'_i}(t) \equiv d_{e'_i}$ for all $i > 0$, and define $d_{e_T}(t) = b_0(d_{e'_0}(t), d_{e'_1}, \dots, d_{e_{n'_1-2}})$, for b_0 from Proposition 1.12 of [5].

Note that for the tuple $\mathbf{d}'_1 = (\mathbf{d}_{\mathcal{E}'_1}, \mathbf{d}_{\mathcal{F}'_1})$ that takes its entries from \mathbf{d} , where \mathcal{E}'_1 is the edge set of T'_1 , we have $P_{v'_T}(\mathbf{d}'_1) = P_{v'_T}(\mathbf{d}) \in \mathcal{BC}_{n'_T}$. If we take v'_T as the root vertex of T'_1 it satisfies the hypotheses of Lemma 2.1, so by that result $\mathbf{d}_{\mathcal{F}'_1} \in \mathcal{BC}_{n'_1}$. In particular, $d_{e_T} = b_0(d_{e'_0}, \dots, d_{e_{n'_1-1}})$ by [5, Prop. 1.12], so $\mathbf{d}_{\mathcal{F}'_1}(0) = \mathbf{d}_{\mathcal{F}'_1}$. We now apply Lemma 2.2 to

produce a deformation $\mathbf{d}_{\mathcal{E}'_1}(t)$ of $\mathbf{d}_{\mathcal{E}'_1}$ such that $\mathbf{d}'_1(t) = (\mathbf{d}_{\mathcal{E}'_1}(t), \mathbf{d}_{\mathcal{F}'_1}(t))$ satisfies the hypotheses of Lemma 2.1.

On the other side we define $T_1 = T_0 - (T' \cup \text{int}(e_T))$, call \mathcal{F}_1 its frontier in V and let $\mathbf{d}_{\mathcal{F}_1}$ be the tuple that takes its entries from \mathbf{d} . We define a deformation $\mathbf{d}_{\mathcal{F}_1}(t)$ of $\mathbf{d}_{\mathcal{F}_1}$ by taking $d_{e_T}(t)$ as prescribed above and $d_f(t) \equiv d_f$ for each other edge $f \in \mathcal{F}_1$. Let \mathcal{E}_1 be the edge set of T_1 and v_T its root vertex. The tuple $\mathbf{d}_1 = (\mathbf{d}_{\mathcal{E}_1}, \mathbf{d}_{\mathcal{F}_1})$ taking entries from \mathbf{d} then satisfies the hypotheses of Lemma 2.1, inheriting properties (1) and (2) there from T (recalling for (2) that $e_T \in \mathcal{F}_1$) and (3) from T_0 . We thus apply Lemma 2.2 to produce $\mathbf{d}_{\mathcal{E}_1}(t)$ such that $\mathbf{d}_1(t) = (\mathbf{d}_{\mathcal{E}_1}(t), \mathbf{d}_{\mathcal{F}_1}(t))$ satisfies the hypotheses of Lemma 2.1.

Note that $(\mathcal{E}_1 \cup \mathcal{F}_1) \cap (\mathcal{E}'_1 \cup \mathcal{F}'_1) = \{e_T\}$, and we have defined $\mathbf{d}_1(t)$ and $\mathbf{d}'_1(t)$ so that their entries corresponding to e_T agree. We now define $\mathbf{d}(t) = (\mathbf{d}_{\mathcal{E}}(t), \mathbf{d}_{\mathcal{F}}(t))$ by taking each $d_e(t)$ from $\mathbf{d}_1(t)$ if $e \in \mathcal{E}_1 \cup \mathcal{F}_1$, from $\mathbf{d}'_1(t)$ if $e \in \mathcal{E}'_1 \cup \mathcal{F}'_1$, and otherwise letting $d_e(t) \equiv d_e$. Note that the only entries of $\mathbf{d}_{\mathcal{F}_1}(t)$ and $\mathbf{d}_{\mathcal{F}'_1}(t)$ that change with t correspond to e_T and e'_0 , each of which lies in \mathcal{E} , so $\mathbf{d}_{\mathcal{F}}(t) \equiv \mathbf{d}_{\mathcal{F}}$. The proof of the Proposition will be completed by showing first that $\mathbf{d}(t) \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$, then that $D_T(\mathbf{d}(t))$ is decreasing, for all small enough t .

To show that $\mathbf{d}(t) \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$ we check the criteria of Definition 1.1, beginning with (1). For all vertices v outside T_0 , $P_v(\mathbf{d}(t)) \equiv P_v(\mathbf{d}) \in \mathcal{AC}_{n_v} - \mathcal{C}_{n_v}$ has maximum entry d_{e_v} by hypothesis. For $v \in T_0^{(0)} - \{v_T, v'_T\}$, Lemma 2.4 asserts that $P_v(\mathbf{d}) \in \mathcal{AC}_{n_v} - (\mathcal{C}_{n_v} \cup \mathcal{BC}_{n_v})$, so $P_v(\mathbf{d}(t))$ remains here for small $t > 0$ since $\mathcal{AC}_{n_v} - (\mathcal{C}_{n_v} \cup \mathcal{BC}_{n_v})$ is open. And we chose $\mathbf{d}_{\mathcal{F}'_1}(t) \in \mathcal{BC}_{n'}$ for all t , and $\mathbf{d}_{\mathcal{E}'_1}(t)$ so that $\mathbf{d}'_1 = (\mathbf{d}_{\mathcal{E}'_1}(t), \mathbf{d}_{\mathcal{F}'_1}(t))$ satisfies the hypotheses of Lemma 2.1, so that result implies that $P_{v'_T}(\mathbf{d}(t)) \in \mathcal{BC}_{n'_T}$ for all t . For each $v \in T^{(0)} - \{v_T\}$, d_{e_v} is the unique maximal entry of $P_v(\mathbf{d})$ since it is not in \mathcal{C}_{n_v} (this follows from [5, Prop. 2.2]), so $d_{e_v}(t)$ remains the maximal entry for small $t > 0$.

For criterion (2) of Definition 1.1 we note first that $\mathbf{d}_{\mathcal{F}_1} \in \mathcal{BC}_{n_1}$ by Lemma 2.1, where $n_1 = |\mathcal{F}_1|$, since $P_{v_T}(\mathbf{d}) = P_{v_T}(\mathbf{d}_1) \in \mathcal{BC}_{n_T}$ by hypothesis. If we enumerate \mathcal{F}_1 as $\{e_T, e_1, \dots, e_{n_1-1}\}$ then $d_{e_T} > d_{e_i}$ for all i , again by Lemma 2.1, since d_{e_T} is by hypothesis maximal among the d_e for e containing v_T . It therefore follows from Proposition 1.12 of [5] that $d_{e_T} = b_0(d_{e_1}, \dots, d_{e_{n_1-1}})$ for b_0 as defined there. For $t \geq 0$, our definition of $d_{e_T}(t)$ and the chain rule give $\frac{d}{dt}d_{e_T}(t) = -\frac{\partial b_0}{\partial d_{e'_0}}$. That this is negative follows from:

Lemma 2.6. *The function $b_0: (\mathbb{R}^+)^{n-1} \rightarrow \mathbb{R}^+$ defined in Proposition 1.12 of [5] satisfies $0 < \frac{\partial}{\partial d_i}b_0(d_1, \dots, d_{n-1}) < 1$ for each i .*

We will prove Lemma 2.6 after finishing the current proof. It implies that for all $t > 0$:

$$d_{e_T}(t) < b_0(d_{e_1}(t), \dots, d_{e_{n_1-1}}(t)) \equiv b_0(d_{e_1}, \dots, d_{e_{n_1-1}})$$

It follows that $\mathbf{d}_{\mathcal{F}_1}(t) \in \mathcal{C}_{n_1}$ (see [5, Cor. 4.10]), and hence that $P_{v_T}(\mathbf{d}(t)) = P_{v_T}(\mathbf{d}_1(t)) \in \mathcal{C}_{n_T}$, since we constructed $\mathbf{d}_1(t)$ to satisfy the hypotheses of Lemma 2.1.

For criterion (3) of Definition 1.1 we note that by construction and Lemma 2.2, $J(P_v(\mathbf{d}(t))) = J(\mathbf{d}_{\mathcal{F}_1}(t))$ for all vertices v of T_1 and $J(P_v(\mathbf{d}(t))) = J(\mathbf{d}_{\mathcal{F}'_1}(t))$ for v in T'_1 . For any vertex v of T that lies outside T_0 , $P_v(\mathbf{d}(t)) \equiv P_v(\mathbf{d})$ for all t , and the strict inequality $J(P_v(\mathbf{d})) < J(P_w(\mathbf{d}))$

for any $w \in T_0^{(0)}$ is preserved for small t . The only vertex of T_0 that does not lie in T_1 or T'_1 is v'_0 , so to check (3) we must only establish that $J(\mathbf{d}_{\mathcal{F}_1}(t)) \geq J(\mathbf{d}_{\mathcal{F}'_1}(t)) \geq J(P_{v'_0}(\mathbf{d}(t)))$ for all small $t > 0$.

We first address $J(\mathbf{d}_{\mathcal{F}_1}(t))$. Applying the chain rule gives:

$$(2.5.3) \quad \frac{d}{dt}J(\mathbf{d}_{\mathcal{F}_1}(t)) = \frac{\partial J}{\partial d_{e_T}} \frac{d}{dt}d_{e_T}(t) + \sum_{i=0}^{n'_1-2} \frac{\partial J}{\partial d_{e_i}} \frac{d}{dt}d_{e_i}(t)$$

By construction, $\frac{d}{dt}d_{e_i}(t) \equiv 0$ for all $i \geq 1$, so the quantity inside the summation above vanishes. Proposition 1.14 of [5] implies that if $\mathbf{d} = (d_0, \dots, d_{n-1})$ lies in \mathcal{C}_n then $0 < \frac{\partial J}{\partial d_i}(\mathbf{d}) < \frac{1}{2}$ for all i , and if $\mathbf{d} \in \mathcal{BC}_n$ then $\frac{\partial J}{\partial d_i}(\mathbf{d}) = \frac{1}{2}$ if d_i is the largest entry. Applying this result and the observation above that $\mathbf{d}_{\mathcal{F}_1}(0) \in \mathcal{BC}_{n_1}$ with largest entry d_{e_T} , and $\mathbf{d}_{\mathcal{F}_1}(t) \in \mathcal{C}_{n_1}$ for small $t > 0$, gives that $\frac{d}{dt}J(\mathbf{d}_{\mathcal{F}_1}(0)) = \frac{1}{2} \frac{d}{dt}d_{e_T}(0)$ and $\frac{1}{2} \frac{d}{dt}d_{e_T}(t) < \frac{d}{dt}J(\mathbf{d}_{\mathcal{F}_1}(t)) < 0$ for $t > 0$. (Here recall from above Lemma 2.6 that $\frac{d}{dt}d_{e_T}(t) < 0$ for all t .)

Applying the chain rule to $\frac{d}{dt}J(\mathbf{d}_{\mathcal{F}'_1}(t))$ in the same way as in (2.5.3) gives $\frac{d}{dt}J(\mathbf{d}_{\mathcal{F}_1}(t)) = \frac{\partial J}{\partial d_{e_T}} \frac{d}{dt}d_{e_T}(t) + \frac{\partial J}{\partial d_{e'_0}} \frac{d}{dt}d_{e'_0}(t)$. We chose $d_{e_T}(t)$ so that $\mathbf{d}_{\mathcal{F}'_1}(t) \in \mathcal{BC}_{n'_1}$ for all t , so [5, Prop. 1.14] implies that $\frac{\partial J}{\partial d_{e'_0}} \equiv 0$ and $\frac{\partial J}{\partial d_{e_T}} \equiv \frac{1}{2}$. (These equalities follow from the inequalities recorded in that result by continuity of the partial derivatives of J , recalling from [5] that \mathcal{BC}_n is the frontier of \mathcal{C}_n in \mathcal{AC}_n for any n , and from above that $d_{e_T}(t)$ is the largest entry of $\mathbf{d}_{\mathcal{F}_1}(t)$ for small $t \geq 0$.) It follows that $\frac{d}{dt}J(\mathbf{d}_{\mathcal{F}'_1}(t)) = \frac{1}{2} \frac{d}{dt}d_{e_T}(t)$ for all t , hence that $J(\mathbf{d}_{\mathcal{F}_1}(t)) \geq J(\mathbf{d}_{\mathcal{F}'_1}(t))$ for all $t \geq 0$.

We finally compute $\frac{d}{dt}J(P_{v'_0}(\mathbf{d}(t)))$ as in (2.5.3), yielding $\frac{\partial J}{\partial d_{e'_0}} \frac{d}{dt}d_{e'_0}(t) = -\frac{\partial J}{\partial d_{e'_0}}$. We established above that $P_{v'_0}(\mathbf{d}(t))$ remains in $\mathcal{AC}_{n_{v'_0}} - (\mathcal{C}_{n_{v'_0}} \cup \mathcal{BC}_{n_{v'_0}})$ for all small $t \geq 0$, so [5, Prop. 1.14] implies that $\frac{\partial J}{\partial d_{e'_0}} > \frac{1}{2}$ for all such t . Therefore $\frac{d}{dt}J(P_{v'_0}(\mathbf{d}(t))) < -\frac{1}{2}$ for all t . But Lemma 2.6 implies that $\frac{d}{dt}d_{e_T}(t) = -\frac{\partial b_0}{\partial d_{e'_0}} > -1$, so $\frac{d}{dt}J(\mathbf{d}_{\mathcal{F}'_1}(t)) > -\frac{1}{2}$ for all $t \geq 0$. It follows that $J(\mathbf{d}_{\mathcal{F}_1}(t)) \geq J(P_{v'_0}(\mathbf{d}(t)))$ for all $t \geq 0$ as claimed, and hence that $\mathbf{d}(t) \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$ for such t .

To show that $D_T(\mathbf{d}(t))$ is decreasing we use the following consequence of our construction and Lemma 2.1:

$$D_T(\mathbf{d}(t)) = D_0(P_{v'_0}(\mathbf{d}(t))) + D_0(\mathbf{d}'_1(t)) + D_0(\mathbf{d}_1(t)) + \sum_{v \in T^{(0)} - T_0^{(0)}} D_0(P_v(\mathbf{d}(t)))$$

We now apply the chain rule and [5, Prop. 2.3] to compute $\frac{d}{dt}D_T(\mathbf{d}(t))$, yielding:

$$(2.5.4) \quad \begin{aligned} & - \left[\sqrt{\frac{1}{\cosh^2(d_{e'_0}(t)/2)} - \frac{1}{\cosh^2 J(\mathbf{d}_{\mathcal{F}'_1}(t))}} - \sqrt{\frac{1}{\cosh^2(d_{e'_0}(t)/2)} - \frac{1}{\cosh^2 J(P_{v'_0}(\mathbf{d}(t)))}} \right] \\ & + \frac{d}{dt}d_{e_T}(t) \sqrt{\frac{1}{\cosh^2(d_{e_T}(t)/2)} - \frac{1}{\cosh^2 J(\mathbf{d}_{\mathcal{F}_1}(t))}} \end{aligned}$$

Here we are using the fact that only $d_{e'_0}(t)$ and $d_{e_T}(t)$ are non-constant among all entries of $\mathbf{d}(t)$; that $d_{e'_0}(t)$ is an entry of $P_{v'_0}(\mathbf{d}(t))$ and $\mathbf{d}_{\mathcal{F}'_1}(t)$, and that $\frac{d}{dt}d_{e'_0}(t) = -1$. This yields the top line above. For the bottom line we use that $d_{e_T}(t)$ is an entry of $\mathbf{d}_{\mathcal{F}'_1}(t)$ and $\mathbf{d}_{\mathcal{F}_1}(t)$, and that $J(\mathbf{d}_{\mathcal{F}'_1}(t)) \equiv d_{e_T}/2$, which implies that $\frac{\partial}{\partial d_{e_T}}D_0(\mathbf{d}_{\mathcal{F}'_1}(t)) \equiv 0$.

The derivative recorded above is 0 at $t = 0$, since there all circumcircle radii are equal to $d_{e_T}/2$. But it is negative for small $t > 0$, since we showed above that $J(\mathbf{d}_{\mathcal{F}'_1}(t)) > J(P_{v'_0}(\mathbf{d}(t)))$ and $\frac{d}{dt}d_{e_T}(t) < 0$ for such t . It follows that $D_T(\mathbf{d}(t))$ decreases with t , and we have proved that if $\mathbf{d} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$ is a local minimizer for $\mathbf{d} \mapsto D_T(\mathbf{d})$ such that $J(P_v(\mathbf{d})) = J(P_w(\mathbf{d}))$ for some $v \in T^{(0)}$ and $w \in v - 1$ then $v_T \in T_0$, $P_{v_T}(\mathbf{d}) \in \mathcal{BC}_{n_T}$, and either $e_T \in \mathcal{F}$ or its other endpoint v'_T lies on the boundary of T_0 .

We have proved the third bulleted assertion of Proposition 2.3. It remains to show for \mathbf{d} as above that if $e_T \in \mathcal{E}$ then $P_v(\mathbf{d}) \in \mathcal{BC}_{n_v}$ for each $v \in v'_T - 1$. Suppose not; ie. that there exists $v'_0 \in v'_T - 1$ such that $P_{v'_0}(\mathbf{d}) \in \mathcal{AC}_{n_{v'_0}} - (\mathcal{C}_{n_{v'_0}} \cup \mathcal{BC}_{n_{v'_0}})$. Let e'_0 denote $e_{v'_0}$, and take $T'_1 = \{v'_T\}$ and $T_1 = T_0 - (T'_1 \cup \text{int}(e_T))$. The frontier $\mathbf{d}_{\mathcal{F}'_1}$ of T'_1 in V is the set of edges containing v'_T , and we take $\mathbf{d}_{\mathcal{F}'_1} = P_{v'_T}(\mathbf{d})$. Letting \mathcal{E}_1 and \mathcal{F}_1 denote the edge set and frontier in V of T_1 , respectively, we define $\mathbf{d}_{\mathcal{F}'_1}(t)$, $\mathbf{d}_{\mathcal{F}_1}(t)$, $\mathbf{d}_{\mathcal{E}_1}(t)$ and $\mathbf{d}(t)$ exactly as in the previous case (note that here $\mathcal{E}'_1 = \emptyset$). That is, we let $d_{e'_0}(t) = d_{e'_0} - t$; choose $d_{e_T}(t)$ as before so that $\mathbf{d}_{\mathcal{F}'_1}(t) \in \mathcal{BC}_{n'_1}(t)$ for all t , where $n'_1 = |\mathcal{F}'_1|$; let Lemma 2.2 determine $\mathbf{d}_{\mathcal{E}_1}(t)$; and let $d_e(t) \equiv d_e$ for all $e \in (\mathcal{E} \cup \mathcal{F}) - (\mathcal{F}'_1 \cup \mathcal{E}_1 \cup \mathcal{F}_1)$.

The same argument as in the previous case now shows that $\mathbf{d}(t) \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$ for small enough $t > 0$, and $D_T(\mathbf{d}(t))$ decreases in t , with only a couple slight modifications. We first note that it is still true that $J(P_{v'_0}(\mathbf{d})) > J(P_v(\mathbf{d}))$ for all $v \in v'_0 - 1$, since we have already showed that if not then v_T is in the maximal subtree containing v'_0 with all vertices v satisfying $J(P_v(\mathbf{d})) = J(P_{v'_0}(\mathbf{d}))$. This fact is necessary for showing that $\mathbf{d}(t) \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$ for small enough $t > 0$. And the computation of $\frac{d}{dt}D_T(\mathbf{d}(t))$ is identical, but in this case the derivative is negative at $t = 0$ since the fact that v'_0 is not in T_0 implies that $J(P_{v'_0}(\mathbf{d})) < J(\mathbf{d}_{\mathcal{F}'_1})$ (in the previous case equality held). But this only helps us, and the result follows. \square

Lemma 2.6. *The function $b_0: (\mathbb{R}^+)^{n-1} \rightarrow \mathbb{R}^+$ defined in Proposition 1.12 of [5] satisfies $0 < \frac{\partial}{\partial d_i}b_0(d_1, \dots, d_{n-1}) < 1$ for each i .*

Proof of Lemma 2.6. The proof of [5, Prop. 1.12] establishes the inequality $\frac{\partial b_0}{\partial d_i} > 0$. There the first of the following equations is showed:

$$\frac{\partial}{\partial d_i}b_0(d_1, \dots, d_{n-1}) = -\frac{\frac{\partial \theta}{\partial d}(d_i, b_0/2)}{\sum_{j=1}^{n-1} \frac{\partial \theta}{\partial J}(d_i, b_0/2)} = \frac{\cosh(d_i/2) \sinh(b_0/2)}{\cosh(b_0/2) \sum_j \sinh(d_j/2)}$$

Here $\theta(d, J)$ is the function described in Lemma 1.4 of [5], that measures the angle of an isosceles triangle with two sides of length J and one of length d at its vertex opposite the side of length d . The latter equation above follows by simply computing partial derivatives.

We recall that $b_0 > d_i$ for each i , by [5, Prop 1.12]. The result now follows by observing that $\sum_j \sinh(d_j/2) > \sinh(b_0/2)$, since by [5, Prop. 1.11], $\mathcal{BC}_n \subset \mathcal{AC}_n$. \square

Corollary 2.7. *Let $T \subset V$ be a compact rooted tree with root vertex $\{v_T\}$, frontier $\mathcal{F} = \{f_0, \dots, f_{n-1}\}$, and edge set \mathcal{E} . For $\mathbf{d}_{\mathcal{F}} = (d_{f_0}, \dots, d_{f_{n-1}}) \in (\mathbb{R}^+)^{\mathcal{F}}$ such that $\overline{Ad}(\mathbf{d}_{\mathcal{F}}) \neq \emptyset$, at a point $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}}) \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$ which is a local minimum of the map $\mathbf{d} \mapsto D_T(\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$, either $P_v(\mathbf{d}) \in \mathcal{BC}_{n_v}$ for each $v \in T^{(0)} - \{v_T\}$ or the following hold.*

$P_{v_T}(\mathbf{d}) \in \mathcal{BC}_{n_T}$, where v_T has valence n_T in V , and for the edge $e_T \in \mathcal{E} \cup \mathcal{F}$ containing v_T such that d_{e_T} is maximal among all such edges, either $e_T \in \mathcal{F}$ or e_T is an edge of the maximal subtree T_0 containing v_T with the property that $J(P_v(\mathbf{d})) = J(P_{v_T}(\mathbf{d}))$ for all $v \in T_0^{(0)}$, with its other endpoint v'_T on the boundary of T_0 . Moreover, for every vertex v of T such that $P_v(\mathbf{d}) \notin \mathcal{BC}_{n_v}$, e_v has its other endpoint in T_0 . In the case that $e_T \in \mathcal{E}$, there is no such $v \in v'_T - 1$.

Proof. Proposition 3.23 of [4] asserts that any local minimizer \mathbf{d} for $\mathbf{d} \mapsto D_T(\mathbf{d})$ on $\overline{Ad}(\mathbf{d}_{\mathcal{F}})$ satisfies one of three criteria that it lists. (The result is only stated there for absolute minimizers, but inspecting its proof shows that it holds for local minimizers.) Criterion (1) is the condition that $P_v(\mathbf{d}) \in \mathcal{BC}_{n_v}$ for all $v \in T^{(0)} - \{v_T\}$. We thus suppose now that \mathbf{d} is a local minimum that does not satisfy (1).

Criterion (2) of [4, Prop. 3.23] is that $P_{v_T} \in \mathcal{BC}_{n_T}$, and criterion (3) is that $J(P_v(\mathbf{d})) = J(P_w(\mathbf{d}))$ for some $v \in T^{(0)}$ and $w \in v - 1$. But Proposition 2.3 implies that if \mathbf{d} satisfies criterion (3) then it also satisfies (2). In fact, defining T_0 as we have here, Proposition 2.3 implies that T_0 contains every $v \in T^{(0)}$ such that $J_v(\mathbf{d}) = J_w(\mathbf{d})$ for some $w \in T^{(0)}$ that is joined to v by an edge of T ; hence that $J_v(\mathbf{d}) = J_{v_T}(\mathbf{d})$ for all such v .

For every vertex v of T outside T_0 , it follows that the inequality of Definition 1.1(3) is strict; ie. that $J(P_v(\mathbf{d})) > J(P_w(\mathbf{d}))$ for all $w \in v - 1$, and moreover, that $J(P_v(\mathbf{d})) < J(P_{v'}(\mathbf{d}))$ for such v , where v' is the other endpoint of e_v . If $P_v(\mathbf{d}) \notin \mathcal{BC}_{n_v}$ and $v' \notin T_0$, then we claim that d_{e_v} can be decreased slightly, keeping all other entries of \mathbf{d} constant, to produce $\mathbf{d}' \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$, with $D_T(\mathbf{d}') < D_T(\mathbf{d})$.

The key effects of the deformation are that $J(P_v(\mathbf{d}')) < J(P_v(\mathbf{d}))$ and $J(P_{v'}(\mathbf{d}')) > J(P_{v'}(\mathbf{d}))$, whereas $J(P_w(\mathbf{d}')) = J(P_w(\mathbf{d}))$ for all other $w \in T^{(0)}$. But since we assumed that $v' \notin T_0$, if w is the other endpoint of $e_{v'}$ then $J(P_w(\mathbf{d})) > J(P_{v'}(\mathbf{d}))$, so this inequality is preserved by choosing d'_{e_v} near enough to d_{e_v} . Similarly, if $w \in v - 1$ the strict inequality $J(P_w(\mathbf{d})) < J(P_v(\mathbf{d}))$ is preserved upon choosing d'_{e_v} near enough to d_{e_v} . Arguing as in the next-to-last paragraph of the proof of [4, Prop. 3.23] shows that $P_{v'}(\mathbf{d}')$ remains in $\mathcal{AC}_{n_{v'}} - \mathcal{C}_{n_{v'}}$, and since $P_v(\mathbf{d})$ lies in the open set $\mathcal{AC}_{n_v} - (\mathcal{C}_{n_v} \cup \mathcal{BC}_{n_v})$ by hypothesis, it remains there for d'_{e_v} near enough to d_{e_v} . The fact that $D_T(\mathbf{d}') < D_T(\mathbf{d})$ follows as in the second paragraph of the proof of [4, Prop. 3.23], and the claim is proved.

The claim implies for all $v \in T^{(0)}$ such that $P_v(\mathbf{d}) \notin \mathcal{BC}_{n_v}$ that e_v has its other endpoint in T_0 . The remaining properties of T_0 follow from Proposition 2.3. We note that as defined here we could have $T_0 = \{v_T\}$, if \mathbf{d} satisfies criterion (2) of [4, Prop. 3.23] but not (3), but in this case e_T must lie in \mathcal{F} . This is because the fact that $P_{v_T}(\mathbf{d}) \in \mathcal{BC}_{n_T}$ implies that $J(P_{v_T}(\mathbf{d})) = d_{e_T}/2$ [5, Prop. 1.11], so if $e_T \in \mathcal{E}$ then the fact that its other endpoint v'_T must

satisfy $J(P_{v'_T}(\mathbf{d})) \leq J(P_{v_T}(\mathbf{d}))$ (by Definition 1.1(1)) but $J(P_{v'_T}(\mathbf{d})) \geq d_{e_T}/2$ [5, Prop. 1.5] implies that $J(P_{v'_T}(\mathbf{d})) = d_{e_T}/2 = J(P_{v_T}(\mathbf{d}))$. \square

2.2. Allowing $\mathbf{d}_{\mathcal{F}}$ and T to vary. The first main result of this subsection, Proposition 2.8, generalizes and strengthens Proposition 3.30 of [4]. The idea here is to bound $D_T(\mathbf{d})$ below for a fixed compact, rooted tree T with frontier \mathcal{F} , but with $\mathbf{d}_{\mathcal{F}}$ allowed to vary with its entries bounded below by those of some fixed $\mathbf{b}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$. After this we prove Lemma 2.10, which compares minima of $D_T(\mathbf{d})$ for different trees T , then prove Theorem 2.11.

Proposition 2.8. *Suppose $T \subset V$ is a compact, rooted tree with root vertex v_T , edge set \mathcal{E} , and frontier \mathcal{F} . Given $\mathbf{b}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$, define $\mathbf{b}_{\mathcal{E}} \in (\mathbb{R}^+)^{\mathcal{E}}$ by $b_e = b_e(\mathbf{b}_{\mathcal{F}})$ for each $e \in \mathcal{E}$, where $b_e: (\mathbb{R}^+)^{\mathcal{F}} \rightarrow \mathbb{R}^+$ is as in [4, Lemma 3.19]. Enumerating the edges of $\mathcal{E} \cup \mathcal{F}$ containing v_T as e_0, \dots, e_{n_T-1} so that b_{e_0} is maximal, define $m_{e_0} = \min\{b_{e_0}, b_0(b_{e_1}, \dots, b_{e_{n_T-1}})\}$, and take:*

$$B_T(\mathbf{b}_{\mathcal{F}}) = D_0(m_{e_0}, b_{e_1}, \dots, b_{e_{n_T-1}}) + \sum_{v \in T^{(0)} - \{v_T\}} D_0(P_v(\mathbf{b}_{\mathcal{E}}, \mathbf{b}_{\mathcal{F}}))$$

Then for each $\mathbf{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$ such that $d_f \geq b_f$ for each $f \in \mathcal{F} - \{e_0\}$ and $d_{e_0} \geq m_{e_0}$, and each $\mathbf{d} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$, $D_T(\mathbf{d}) \geq B_T(\mathbf{b}_{\mathcal{F}})$.

If $e_0 \notin \mathcal{F}$ then the requirement above on $\mathbf{d}_{\mathcal{F}}$ simply becomes that $d_f \geq b_f$ for all $f \in \mathcal{F}$. In the case that $e_0 \in \mathcal{F}$ we note that the given bound is *a priori* stronger than one which holds for all $\mathbf{d}_{\mathcal{F}}$ with $d_f \geq b_f$ for all $f \in \mathcal{F}$, since $m_{e_0} \leq b_{e_0}$.

The proof proceeds by separately considering the possibilities described in Corollary 2.7 for minimizers of the function $\mathbf{d} \mapsto D_T(\mathbf{d})$. One is handled by the lemma below.

Lemma 2.9. *Suppose $T \subset V$ is a rooted tree with root vertex v_T , edge set \mathcal{E} , and frontier \mathcal{F} , and fix $\mathbf{b}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$. Suppose for $\mathbf{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$ and $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}}) \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$, that*

- (1) $P_v(\mathbf{d}) \in \mathcal{BC}_{n_v}$ for each $v \in T^{(0)} - \{v_T\}$; and
- (2) $d_f \geq b_f$ for each $f \in \mathcal{F} - \{e_0\}$, and $d_{e_0} \geq m_{e_0}$, for e_0 and m_{e_0} as in Proposition 2.8.

Then $D_T(\mathbf{d}) \geq B_T(\mathbf{b}_{\mathcal{F}})$, for B_T as defined in Proposition 2.8.

Proof. For a given \mathbf{d} , hypothesis (1) above and the defining property of the functions b_e from [4, Lemma 3.19] imply for each $e \in \mathcal{E}$ that $d_e = b_e(\mathbf{d}_{\mathcal{F}})$. The monotonicity property of b_e laid out in assertion (3) of [4, L. 3.19] and the hypothesis that $d_f \geq b_f$ for each $f \in \mathcal{F}$ thus together imply that $d_e \geq b_e$ for each $e \in \mathcal{E}$. Corollary 2.4 of [5] now directly gives for each $v \in T^{(0)} - \{v_T\}$ that $D_0(P_v(\mathbf{b}_{\mathcal{E}}, \mathbf{b}_{\mathcal{F}})) \leq D_0(P_v(\mathbf{d}))$.

For e_0, \dots, e_{n_T-1} as in Proposition 2.8, with m_{e_0} as defined there we note that by construction $(m_{e_0}, d_{e_1}, \dots, d_{e_{n_T-1}})$ lies in $\mathcal{C}_{n_T} \cup \mathcal{BC}_{n_T}$. Therefore since $m_{e_0} \leq b_{e_0}$ and $P_{v_T}(\mathbf{d}) \in \mathcal{C}_{n_T} \cup \mathcal{BC}_{n_T}$ by Definition 1.1(2), [5, Cor. 2.4] also implies that $D_0(m_{e_0}, b_{e_1}, \dots, b_{e_{n_T-1}}) \leq D_0(P_{v_T}(\mathbf{d}))$. The Lemma therefore follows from the definitions of $B_T(\mathbf{b}_{\mathcal{F}})$ and $D_T(\mathbf{d}_{\mathcal{F}})$. \square

Proof of Proposition 2.8. For $\mathbf{b}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$ and T as in the Proposition, compute $\mathbf{b}_{\mathcal{E}}$ and m_{e_0} as prescribed there. Now enumerate \mathcal{F} as $\{f_1, \dots, f_n\}$, where $f_n = e_0$ if $e_0 \in \mathcal{F}$. For the set $SAd_T = \{\mathbf{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}} \mid \overline{Ad}(\mathbf{d}_{\mathcal{F}}) \neq \emptyset\}$ defined in Lemma 3.29 of [4], that result implies that the set below is closed in \mathbb{R}^n .

$$SAd_T(\mathbf{b}_{\mathcal{F}}) = SAd_T \cap \{(d_1, \dots, d_n) \in \mathbb{R}^n \mid d_i \geq b_{f_i} \text{ for each } i < n, \text{ and } d_n \geq m_{e_0} \text{ or } b_{f_n}\}$$

Here the inequality $d_n \geq m_{e_0}$ applies if $e_0 \in \mathcal{F}$, hence $f_n = e_0$; otherwise we require $d_n \geq b_{f_n}$.

Since $SAd_T(\mathbf{b}_{\mathcal{F}})$ is closed its intersection with $[0, D]^n$ is compact for any fixed $D > 0$. Let D be large enough that $SAd_T(\mathbf{b}_{\mathcal{F}}) \cap [0, D]^n$ is non-empty. (It is not hard to show that such a D exists.) Then [4, L. 3.29] further implies that the function $\mathbf{d}_{\mathcal{F}} \mapsto \min\{D_T(\mathbf{d}) \mid \mathbf{d} \in \overline{Ad}_T(\mathbf{d}_{\mathcal{F}})\}$ attains a minimum on it, since it asserts that this function is lower-semicontinuous on SAd_T .

Let $\mathbf{d}_{\mathcal{F}}$ be a minimizer for $\min\{D_T(\mathbf{d}) \mid \mathbf{d} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})\}$ on $SAd_T(\mathbf{b}_{\mathcal{F}}) \cap [0, D]^n$, and \mathbf{d} a minimizer for $D_T(\mathbf{d})$ on $\overline{Ad}(\mathbf{d}_{\mathcal{F}})$. We apply Corollary 2.7 to \mathbf{d} , separately treating the different cases it describes. If $P_v(\mathbf{d}) \in \mathcal{BC}_{n_v}$ for all $v \in T^{(0)} - \{v_T\}$ then Lemma 2.9 directly implies the desired bound. So we will assume now that $P_v(\mathbf{d}) \notin \mathcal{BC}_{n_v}$ for some $v \in T^{(0)} - \{v_T\}$, and therefore by Cor. 2.7 that $P_{v_T}(\mathbf{d}) \in \mathcal{BC}_{n_T}$, where v_T has valence n_T as in the Corollary.

Let e_T be the edge containing v_T such that d_{e_T} is maximal among all such edges. We first suppose that $e_T \in \mathcal{F}$. In this case we will deform $\mathbf{d}_{\mathcal{F}}$ within $SAd_T(\mathbf{b}_{\mathcal{F}}) \cap [0, D]^n$ to reduce $D_T(\mathbf{d})$, thereby contradicting our minimality hypothesis. The idea is to reduce d_{e_T} without changing any other entry of $\mathbf{d}_{\mathcal{F}}$, so we first show there is room to do this in $SAd_T(\mathbf{b}_{\mathcal{F}})$. With the edges containing v_T numbered e_0, \dots, e_{n_T-1} so that b_{e_0} is maximal, we have $e_T = e_{i_0}$ for some i_0 . We claim that if $i_0 \neq 0$ then $d_{e_T} > b_{e_T} = b_{e_{i_0}}$, and if $i_0 = 0$ then $d_{e_T} > m_{e_0}$.

(The strict inequalities are the point here, since the hypothesis that $d_{f_i} \geq b_{f_i}$ for each $f_i \in \mathcal{F} - \{e_T\}$ implies that $d_e \geq b_e$ for each $e \in \mathcal{E}$, by assertions (2) and (3) of [4, L. 3.19]. This is not clear directly from that lemma's statement, since we do not necessarily have $d_{e_T} \geq b_{e_T}$ here, but note that the first sentence of its proof asserts for each $v \in T^{(0)} - \{v_T\}$ that $b_{e_v}(\mathbf{d}_{\mathcal{F}})$ is determined by $\mathbf{d}_{\mathcal{F}}$ and the collection of $b_{e_w}(\mathbf{d}_{\mathcal{F}})$ for $w < v$. In fact the proof itself shows more precisely that $b_{e_v}(\mathbf{d}_{\mathcal{F}})$ is determined by the set of d_{f_i} such that $f_i \in \mathcal{F}$ contains v or some $w < v$. Here e_T contains only v_T , and of course $v_T \not\prec v$ for any v .)

First suppose that $i_0 \neq 0$. By the above $d_{e_i} \geq b_{e_i}$ for each $i \neq i_0$, and in particular $d_{e_0} \geq b_{e_0}$. The claim now follows from the fact that since $P_{v_T}(\mathbf{d}) \in \mathcal{BC}_{n_T}$,

$$d_{e_T} = b_0(d_{e_0}, \dots, \widehat{d_{e_{i_0}}}, \dots, d_{e_{n_T-1}}),$$

so recalling [5, Prop. 1.12] we have that it is larger than $d_{e_0} \geq b_{e_0}$.

Now suppose that $i_0 = 0$. In this case we use our assumption that $P_v(\mathbf{d}) \notin \mathcal{BC}_{n_v}$ for some $v \in T^{(0)} - \{v_T\}$, which implies for the initial edge e_v of the arc of T joining v to v_T that

$$d_{e_v} > b_0(d_{e'_1}, \dots, d_{e'_{n_v-1}}) \geq b_{e_v} = b_0(b_{e'_1}, \dots, b_{e'_{n_v-1}}),$$

where e'_1, \dots, e'_{n_v-1} are the other edges of $\mathcal{E} \cup \mathcal{F}$ containing v . It now follows from Lemma 2.6 that strict inequality $d_e > b_e$ holds for each edge e on the arc joining v to v_T , using e_v

as the base case of an inductive argument. This holds in particular for one of the edges e_i containing v_T , for $i > 0$, so since $P_{v_T}(\mathbf{d}) \in \mathcal{BC}_{n_T}$ we again obtain from Lemma 2.6 that

$$d_{e_T} = b_0(d_{e_1}, \dots, d_{e_{n_T-1}}) > b_0(b_{e_1}, \dots, b_{e_{n_T-1}}).$$

It thus follows that $d_{e_T} > m_{e_0}$, and the claim is proved in this case as well.

We now define $\mathbf{d}_{\mathcal{F}}(t)$ by taking $d_{e_T}(t) = d_{e_T} - t$ and $d_{f_i}(t) \equiv d_{f_i}$ for all $f_i \in \mathcal{F} - \{e_T\}$. We will next produce a deformation $\mathbf{d}(t) \in \overline{Ad}(\mathbf{d}_{\mathcal{F}}(t))$ of \mathbf{d} for small $t > 0$, from which (together with the claim) it will follow that $\mathbf{d}_{\mathcal{F}}(t)$ is a deformation of $\mathbf{d}_{\mathcal{F}}$ within $SAd_T(\mathbf{b}_{\mathcal{F}}) \cap [0, D]^n$. We will further observe that $D_T(\mathbf{d}(t))$ decreases with t , thus obtaining a contradiction to the hypothesis that $\mathbf{d}_{\mathcal{F}}$ is a minimizer for $\min\{D_T(\mathbf{d}) \mid \mathbf{d} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})\}$ on $SAd_T(\mathbf{b}_{\mathcal{F}}) \cap [0, D]^n$.

Let T_0 be the maximal subtree of T such that $J(P_v(\mathbf{d})) = J(P_{v_T}(\mathbf{d}))$ for all $v \in T_0^{(0)}$. Define $\mathbf{d}(t)$ by letting $d_e(t) \equiv d_e$ for each edge e of T that does not lie in T_0 , and for $e \subset T_0$ let $d_e(t)$ be determined by Lemma 2.2. The verification that $\mathbf{d}(t) \in \overline{Ad}(\mathbf{d}_{\mathcal{F}}(t))$ parallels the corresponding check for the deformation $\mathbf{d}(t)$ described in the proof of the third bulleted assertion of Proposition 2.3, but it is simpler. In particular, T_0 here plays the role of T_1 there, but there is no T'_1 or v'_0 , and here we have $\frac{d}{dt}d_{e_T}(t) \equiv -1$ rather than $\frac{d}{dt}d_{e_T}(t) = -\frac{\partial b_0}{\partial d_{e'_0}} < 0$ as there (see above the statement of Lemma 2.6). As was the case with $\mathbf{d}_{\mathcal{F}_1}(t)$, we have $\mathbf{d}_{\mathcal{F}_0}(t) \in \mathcal{C}_{n_0}$ for small $t > 0$, where \mathcal{F}_0 is the frontier of T_0 and $n_0 = |\mathcal{F}_0|$. We thus obtain

$$\frac{d}{dt}D_T(\mathbf{d}(t)) = -\sqrt{\frac{1}{\cosh^2(d_{e_T}(t)/2)} - \frac{1}{\cosh^2 J(\mathbf{d}_{\mathcal{F}_0}(t))}} < 0$$

for small $t > 0$. This is the analog of (2.5.4) in the proof of Proposition 2.3, but again it is simpler since there is no T'_1 or v'_0 , and here $e_T \in \mathcal{F}$ instead. It implies that $D_T(\mathbf{d}(t))$ is decreasing as asserted, finishing the case that $P_{v_T}(\mathbf{d}) \in \mathcal{BC}_{n_T}$ and $e_T \in \mathcal{F}$.

We finally address the case that $P_{v_T}(\mathbf{d}) \in \mathcal{BC}_{n_T}$ and $e_T \in \mathcal{E}$. Here we will argue by induction on the number of edges of T , ie. $|\mathcal{E}|$. The base case $T = \{v_T\}$, with $\mathcal{E} = \emptyset$, follows directly from Lemma 2.9, since criterion (1) there holds vacuously (here recall Remark 1.2.) So assume now that T has $k \geq 1$ edges, that the Proposition holds for all trees with fewer than k edges, and that $P_{v_T}(\mathbf{d}) \in \mathcal{BC}_{n_T}$ with $e_T \in \mathcal{E}$.

Let T' be the maximal subtree containing the other endpoint v'_T of e_T but not v_T , and take $T'' = T - (T' \cup \text{int}(e_T))$. We take v'_T as the root vertex of T' and v_T as the root vertex of T'' . Naming the frontiers of T' and T'' as \mathcal{F}' and \mathcal{F}'' , respectively, we have $\mathcal{F}' \cap \mathcal{F}'' = \{e_T\}$ and $\mathcal{F}' \cup \mathcal{F}'' = \mathcal{F} \cup \{e_T\}$. Similarly taking their edge sets to be \mathcal{E}' and \mathcal{E}'' , we have $\mathcal{E}' \cap \mathcal{E}'' = \emptyset$ and $\mathcal{E}' \cup \mathcal{E}'' \cup \{e_T\} = \mathcal{E}$. Thus by the induction hypothesis the Proposition holds for T' and T'' .

Let $\mathbf{d}' = (\mathbf{d}'_{\mathcal{E}'}, \mathbf{d}'_{\mathcal{F}'})$ and $\mathbf{d}'' = (\mathbf{d}''_{\mathcal{E}''}, \mathbf{d}''_{\mathcal{F}''})$ take their entries from \mathbf{d} , so in particular $d'_{e_T} = d''_{e_T} = d_{e_T}$. Then for any vertex v of T' , $P_v(\mathbf{d}') = P_v(\mathbf{d})$, and similarly if $v \in T''$. Therefore \mathbf{d}' lies in $\overline{Ad}(\mathbf{d}'_{\mathcal{F}'})$, and $\mathbf{d}'' \in \overline{Ad}(\mathbf{d}''_{\mathcal{F}''})$: criteria (1) - (3) of Definition 1.1 are directly inherited by T'' from T , and for T' we merely note in addition that $P_{v'_T}(\mathbf{d}') \in \mathcal{BC}_{n'_T}$ by hypothesis, and $J(P_{v'_T}(\mathbf{d}')) = J(P_{v_T}(\mathbf{d})) = d_{e_T}/2$. Note that $D_T(\mathbf{d}) = D_{T'}(\mathbf{d}') + D_{T''}(\mathbf{d}'')$, since each vertex of T lies in exactly one of T' or T'' .

Define $\mathbf{b}'_{\mathcal{F}'}$ by taking $b'_{e_T} = b_{e_T} = b_{e_T}(\mathbf{b}_{\mathcal{F}})$ and pulling the remaining entries from $\mathbf{b}_{\mathcal{F}}$, and define $\mathbf{b}''_{\mathcal{F}''}$ by taking all entries but b''_{e_T} from $\mathbf{b}_{\mathcal{F}}$. To obtain b''_{e_T} we enumerate the edges containing v_T as e_0, \dots, e_{n_T-1} as described in the Proposition, for each i let b_{e_i} be the appropriate entry of $\mathbf{b}_{\mathcal{F}}$ or $\mathbf{b}_{\mathcal{E}}$, and for i_0 such that $e_T = e_{i_0}$ we take

$$b''_{e_T} = b_0(b_{e_0}, \dots, \widehat{b_{e_{i_0}}}, \dots, b_{e_{n_T-1}})$$

We will establish this case of the Proposition with two claims: first, that $B_{T'}(\mathbf{d}'_{\mathcal{F}'}) + B_{T''}(\mathbf{d}''_{\mathcal{F}''}) \geq B_T(\mathbf{d}_{\mathcal{F}})$, and second, that $d'_e \geq b'_e$ for all $e \in \mathcal{F}'$ and $d''_e \geq b''_e$ for all $e \in \mathcal{F}''$. Applying induction and the observation that $D_T(\mathbf{d}) = D_{T'}(\mathbf{d}') + D_{T''}(\mathbf{d}'')$, we therefore conclude the desired bound $D_T(\mathbf{d}) \geq B_T(\mathbf{b}_{\mathcal{F}})$.

Toward the first claim, computing straight from the definitions gives:

$$B_{T'}(\mathbf{d}'_{\mathcal{F}'}) + B_{T''}(\mathbf{d}''_{\mathcal{F}''}) - B_T(\mathbf{d}_{\mathcal{F}}) = D_0(b_{e_0}, \dots, b''_{e_T}, \dots, b_{e_{n_T-1}}) - D_0(m_{e_0}, b_{e_1}, \dots, b_{e_{n_T-1}})$$

The entries from the two inputs to D_0 on the right-hand side above differ only in the e_0 and e_{i_0} positions. If $i_0 = 0$ then it follows directly from their definitions that $m_{e_0} \leq b''_{e_T}$. Otherwise, by its definition in the Proposition we have $m_{e_0} \leq b_{e_0}$. And since b_{e_0} is maximal among the b_{e_i} by definition, $b''_{e_T} > b_{e_0} \geq b_{e_{i_0}}$ by its definition and Proposition 1.12 of [5]. In either case the difference above is positive by [5, Cor. 2.4], yielding the first claim.

For the second claim we note first that by hypothesis $d_e \geq b_e$ for each $e \in \mathcal{F}$, so by definition $d'_e = d_e \geq b_e$ for all $e \in \mathcal{F}' - \{e_T\}$, and similarly for $e \in \mathcal{F}'' - \{e_T\}$. Applying Lemma 3.19 of [4] gives

$$d_e \geq b_e(\mathbf{d}_{\mathcal{F}}) \geq b_e(\mathbf{b}_{\mathcal{F}}) = b_e$$

for each $e \in \mathcal{E}$. (The first and second inequalities above are respectively implied by assertions (2) and (3) there.) It now follows immediately that $d'_{e_T} = d_{e_T} \geq b'_{e_T} = b_{e_T}$, so the claim is proved for $\mathbf{d}_{\mathcal{F}'}$. Since $P_{v_T}(\mathbf{d}) \in \mathcal{C}_{n_T}$, Proposition 1.12 of [5] implies that

$$d_{e_T} = b_0(d_{e_0}, \dots, \widehat{d_{e_{i_0}}}, \dots, d_{e_{n_T-1}}) \geq b_0(b_{e_0}, \dots, \widehat{b_{e_{i_0}}}, \dots, b_{e_{n_T-1}})$$

But the latter quantity is b''_{e_T} , and since $d''_{e_T} = d_{e_T}$ the second claim is also proved for \mathcal{F}'' . \square

Lemma 2.10. *Suppose $T \subset V$ is a compact, rooted tree with root vertex v_T , edge set \mathcal{E} and frontier \mathcal{F} , where each vertex of T has valence at least three in V . There is a compact, rooted tree $T_0 \subset V_0$ with root vertex v_{T_0} , edge set \mathcal{E}_0 and frontier \mathcal{F}_0 , where each vertex of T_0 is trivalent in V_0 , with the following property.*

There is a bijection $q: \mathcal{F} \rightarrow \mathcal{F}_0$ such that for any $\mathbf{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$, the tuple $\mathbf{d}_{\mathcal{F}_0} \in (\mathbb{R}^+)^{\mathcal{F}_0}$ given by relabeling entries of $\mathbf{d}_{\mathcal{F}}$ using q has the property that:

$$\min\{D_T(\mathbf{d}) \mid \mathbf{d} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})\} \geq \min\{D_{T_0}(\mathbf{d}) \mid \mathbf{d} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}_0})\}$$

Proof. Suppose T has k vertices. Each vertex of T has valence at least three in V , so $3k \leq 2|\mathcal{E}| + |\mathcal{F}|$, with equality holding if and only if each vertex of T is trivalent in V . Since T is a tree its Euler characteristic is 1, so we also have $|\mathcal{E}| = k - 1$. Substituting this into the first inequality gives $k \leq |\mathcal{F}| - 2$, with equality if and only if each vertex of T is trivalent

in V . We will prove the Lemma by fixing $n = |\mathcal{F}| \geq 3$ and inducting on $n - k$. The base case $n - k = 2$ holds trivially with $T_0 = T$ and q the identity map.

Let us now take $k < n - 2$ and suppose the Lemma holds for all trees with frontier of order n and more than k vertices. If $T \subset V$ is a compact, rooted tree with root vertex v_T , edge set \mathcal{E} of order k , and frontier \mathcal{F} of order n such that each vertex of T has valence at least three in V then by the first paragraph there is a vertex v of T with valence at least four in V . List the edges in $\mathcal{E} \cup \mathcal{F}$ containing v as e_0, \dots, e_{n_v-1} , where n_v is the valence of v in V , and if $v \neq v_T$ then $e_0 = e_v$ is the initial edge of the arc in T joining v to v_T .

Let T_1 be the maximal subtree of T containing e_0 and e_1 but not e_i for $i > 1$, and let T_2 be the maximal subtree containing the remaining e_i but not e_0 or e_1 . Then $T = T_1 \cup T_2$ and $T_1 \cap T_2 = \{v\}$. We produce a tree T' with $k + 1$ vertices by joining a copy of T_1 to a copy of T_2 by an edge e' that has its endpoints at the respective copies of v in T_1 and T_2 . We produce V' containing T' similarly, by doubling v and joining the resulting copies by e' .

There is a quotient map $p_{e'}: V' \rightarrow V$ that identifies e to a point and takes T' to T . It induces a bijection from $\mathcal{E}' - \{e'\}$ to \mathcal{E} , where \mathcal{E}' is the edge set of T' , and from the frontier \mathcal{F}' of T' in V' to \mathcal{F} . We will refer by q' to refer to the inverse bijections both $\mathcal{F} \rightarrow \mathcal{F}'$ and $\mathcal{E} \rightarrow \mathcal{E}' - \{e'\}$. Given $\mathbf{d}_{\mathcal{F}} = (d_f \mid f \in \mathcal{F}) \in (\mathbb{R}^+)^{\mathcal{F}}$ one produces $\mathbf{d}'_{\mathcal{F}'} \in (\mathbb{R}^+)^{\mathcal{F}'}$ by relabeling: $\mathbf{d}'_{\mathcal{F}'} = (d_{q(f)} \mid f \in \mathcal{F})$. Similarly, given $\mathbf{d}_{\mathcal{E}} \in (\mathbb{R}^+)^{\mathcal{E}}$, relabeling gives all entries of an element $\mathbf{d}'_{\mathcal{E}'} \in (\mathbb{R}^+)^{\mathcal{E}'}$ but one, $d_{e'}$. If $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}}) \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$, we will choose $d_{e'}$ and a root vertex $v_{T'}$ for T' so that the resulting element $\mathbf{d}' = (\mathbf{d}'_{\mathcal{E}'}, \mathbf{d}'_{\mathcal{F}'})$ lies in $\overline{Ad}(\mathbf{d}'_{\mathcal{F}'})$.

With the edges of T containing v enumerated as above let $d_i = d_{e_i}$ for each $i < n_v$, and define $d_{e'} = \ell_{n_v-1,1}(d_0, \dots, d_{n_v-1})$, where $\ell_{i,j}$ is the diagonal-length function described in Corollary 1.15 of [5]. By that result, $d_{e'}$ is the length of the diagonal γ of a cyclic n_v -gon C_v with side length collection (d_0, \dots, d_{n_v-1}) that cuts off the sides with lengths d_0 and d_1 from the others. We now take $\mathbf{d}'_{\mathcal{E}'}$ as suggested in the previous paragraph, with $d_{q(e)} = d_e$ for each in \mathcal{E} and $d_{e'}$ as given here.

Now we assign T' a root vertex $v_{T'}$. Let v_1 be the copy of v that lies in $T_1 \subset T'$, and let v_2 be the other copy of v in T' . If $v \neq v_T$ we let $v_{T'} = p_{e'}^{-1}(v_T)$, a vertex of T' since $p_{e'}$ is injective away from e' . Now suppose $v = v_T$. If the circumcircle center of C_v lies on the side of γ containing the edges of length d_0 and d_1 then we let $v_1 = v_{T'}$; otherwise we let $v_2 = v_{T'}$.

For $\mathbf{d}' = (\mathbf{d}'_{\mathcal{E}'}, \mathbf{d}'_{\mathcal{F}'})$ as prescribed above, we claim that $\mathbf{d}' \in \overline{Ad}(\mathbf{d}'_{\mathcal{F}'})$. In the case that $v \neq v_T$, condition (2) of Definition 1.1 follows immediately by construction. If $v = v_T$ then since $P_v(\mathbf{d}) \in \mathcal{C}_{n_T} \cup \mathcal{BC}_{n_T}$, Proposition 2.2 of [5] implies that the cyclic n_v -gon C_v described above contains its circumcircle center c . The diagonal γ above divides C_v into cyclic n -gons C_{v_1} and C_{v_2} with respective side length collections $P_{v_1}(\mathbf{d}')$ and $P_{v_2}(\mathbf{d}')$, where $P_{v_1}(\mathbf{d}') = (d_0, d_1, d_{e'})$ and $P_{v_2}(\mathbf{d}') = (d_{e'}, d_2, \dots, d_{n_v-1})$. We chose to label v_1 or v_2 as $v_{T'}$ according to which of C_{v_1} or C_{v_2} contains c , so again by [5, Prop. 2.2] we have $P_{v_{T'}}(\mathbf{d}') \in \mathcal{C}_{n_{T'}} \cup \mathcal{BC}_{n_{T'}}$.

By construction $J(P_{v_1}(\mathbf{d}')) = J(P_{v_2}(\mathbf{d}')) = J(P_v(\mathbf{d}))$, and condition (3) of Definition 1.1 follows. Condition (1) of Definition 1.1 follows for any vertex v' of T' outside e' from the fact that $P_{v'}(\mathbf{d}') = P_{p_{e'}(v')}(\mathbf{d})$. We now check it for v_1 and v_2 .

If $v \neq v_T$ then $P_v(\mathbf{d}) = (d_0, \dots, d_{n_v-1}) \in \mathcal{AC}_{n_v} - \mathcal{C}_{n_v}$ has largest entry d_0 , since we enumerated the edges containing v as e_0, \dots, e_{n_v-1} so that $e_0 = e_v$. Then by [5, Prop. 2.2] the side of C_v with length d_0 separates C_v from its circumcircle center c . It therefore also separates C_{v_1} from c , and γ separates C_{v_2} from c , which is their shared circumcircle center. Thus again by [5, Prop. 2.2], $P_{v_1}(\mathbf{d}') \in \mathcal{AC}_3 - \mathcal{C}_3$ has largest entry d_0 and $P_{v_2}(\mathbf{d}') \in \mathcal{AC}_{n_v-2} - \mathcal{C}_{n_v-2}$ has largest entry $d_{e'}$. Since e_v is the initial edge of the arc joining v to v_T in T , this arc lies in T_1 and joins v_1 to $v_{T'} = v_T$ in $T_1 \subset T'$. Its initial edge is still e_v , and hence e' is the initial edge of the arc in T' joining v_2 to $v_{T'}$.

If $v = v_T$ and $v_2 = v_{T'}$ then γ separates C_{v_1} from c , so again Proposition 2.2 of [5] implies that $d_{e'}$ is the largest entry of $P_{v_1}(\mathbf{d}') \in \mathcal{AC}_3 - \mathcal{C}_3$. Since e' is the arc joining v_1 to $v_2 = v_{T'}$, Definition 1.1(1) follows in this case. The case that $v_1 = v_{T'}$ is completely analogous, and we have proven the claim that $\mathbf{d}' \in \overline{Ad}(\mathbf{d}_{\mathcal{F}'})$.

Our construction of T' and \mathbf{d}' has reverse-engineered the hypotheses of Lemma 3.28 of [4], since $J(P_{v_1}(\mathbf{d}')) = J(P_{v_2}(\mathbf{d}'))$ and T is obtained from T' by crushing e' to a point. In the notation of that result and Definition 3.26 there, $T = T'_{e'}$ and $\mathbf{d} = \mathbf{d}'_{e'}$. Therefore that Lemma gives $D_T(\mathbf{d}) = D_{T'}(\mathbf{d}')$. Choosing \mathbf{d} as a minimizer for D_T over $\overline{Ad}(\mathbf{d}_{\mathcal{F}})$ gives:

$$\min\{D_T(\mathbf{d}) \mid \mathbf{d} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})\} \geq \min\{D_{T'}(\mathbf{d}') \mid \mathbf{d}' \in \overline{Ad}(\mathbf{d}'_{\mathcal{F}'})\}$$

Now applying the induction hypothesis to T' , which has one more vertex than T , we conclude that the Lemma holds for T . Thus by induction the Lemma holds for all trees with frontier of order $n \geq 3$. But n is arbitrary, so the Lemma holds. \square

Theorem 2.11. *Let C be a compact two-cell of the centered dual complex of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$ such that for some $\mathbf{b} = (b_1, \dots, b_n) \in (\mathbb{R}^+)^n$ and enumeration of the edges of C , the i^{th} edge has length at least b_i for each i . Then $\text{area}(C) \geq \min\{B_T(\sigma(\mathbf{b})) \mid T \in \mathcal{T}_n, \sigma \in S_n\}$, where B_T is the area bounding function defined in Proposition 2.8, S_n is the symmetric group on n letters, $\sigma \in S_n$ acts on \mathbf{b} by permutation of entries, and \mathcal{T}_n is the collection of compact, rooted trees T with frontier \mathcal{F} of order n and each vertex trivalent in $V = T \cup \bigcup_{f \in \mathcal{F}} f$.*

Proof. Let $T \subset V$ be the dual tree to C (recall Definition 2.11 of [4]), where V is the Voronoi tessellation's one-skeleton, and enumerate the frontier \mathcal{F} of T as $\{f_1, \dots, f_n\}$ so that the edge of C dual to f_i has length at least b_i for each i . Let d_i be the length of this edge, and let $\mathbf{d}_{\mathcal{F}} = (d_1, \dots, d_n) \in (\mathbb{R}^+)^n$. Taking $\mathbf{d}_{\mathcal{E}}$ to be the tuple of lengths of Delaunay edges dual to edges of T , Lemma 3.14 of [4] implies that $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$ lies in $Ad(\mathbf{d}_{\mathcal{F}}) \subset \overline{Ad}(\mathbf{d}_{\mathcal{F}})$, and the area of C is $D_T(\mathbf{d})$.

By Lemma 2.10 there is a tree $T_0 \subset V_0$ with frontier \mathcal{F}_0 bijective to \mathcal{F} , such that each vertex of T_0 is trivalent in V_0 , with the property that relabeling the entries of $\mathbf{d}_{\mathcal{F}}$ using the bijection $\mathcal{F} \rightarrow \mathcal{F}_0$ yields $\mathbf{d}_{\mathcal{F}_0}$ satisfying:

$$D_T(\mathbf{d}) \geq \min\{D_{T_0}(\mathbf{d}_0) \mid \mathbf{d}_0 \in \overline{Ad}(\mathbf{d}_{\mathcal{F}_0})\}$$

By Proposition 2.8, this quantity in turn is bounded below by $B_{T_0}(\mathbf{b}_{\mathcal{F}_0})$, where $\mathbf{b}_{\mathcal{F}_0}$ is obtained from $\mathbf{b}_{\mathcal{F}}$ by relabeling in the same way. The Theorem follows. \square

Remark 2.12. The area function D_0 is symmetric in its inputs [5, Prop. 2.3], and the semicyclic radius function b_0 is too [5, Prop. 1.12]. Using these facts it is not hard to show that for any \mathbf{b} and tree T , if edges f_1 and f_2 of \mathcal{F} terminate at the same vertex of T then for the transposition τ that swaps the corresponding entries b_{f_1} and b_{f_2} of \mathbf{b} , $b_e(\mathbf{b}) = b_e(\tau(\mathbf{b}))$ for each $e \in \mathcal{E}$ and $B_T(\mathbf{b}) = B_T(\tau(\mathbf{b}))$. So in computing $\min\{B_T(\mathbf{b})\}$ above, for each tree T it is only necessary to test one representative of each left coset of the subgroup S_T (isomorphic to a direct sum of \mathbb{Z}_2 's) of S_n generated by such swaps.

Furthermore, an automorphism f of (T, v_T) has an induced action on \mathbf{b} which is well-defined up to the action of S_T , where the edges of \mathcal{F} that terminate at v are taken to those that terminate at $f(v)$ for each $v \in T^{(0)}$ and the corresponding entries of \mathbf{b} go along for the ride. One can show again that $b_e(\mathbf{b}) = b_e(f(\mathbf{b}))$ for each $e \in \mathcal{E}$, and $B_T(\mathbf{b}) = B_T(f(\mathbf{b}))$. Thus for each tree T it is in fact only necessary to test $B_T(\sigma(\mathbf{b}))$ for representatives σ of each orbit of the action of the automorphism group of (T, v_T) on S_n/S_T .

3. PRACTICE

The main goal of this section is to describe the Python module *minimizer.py* and data file *forest.txt*, which together give us the ability to obtain the bounds of Theorem 2.11 for arbitrary n -tuples, $n \leq 9$, using a computer. First, in subsection 3.1 we record some existing explicit formulas for geometric measurements of cyclic polygons. We use these to give a completely explicit statement of Theorem 2.11, in Corollary 3.4. Then in subsection 3.2 we describe *forest.txt* and the components of *minimizer.py*.

Here is a brief explanation of how to use *minimizer* to compute the bounds of Theorem 2.11. After downloading *minimizer* and *forest* you must first replace “yourpath” on line 66 of *minimizer.py*, with your path to *forest.txt*. Then in a Python interpreter, import *minimizer.py* and run *minimizer.minimize()* on the desired tuples. Here is a sample series of commands at the Python command prompt, to get it up and running:

```
>>> import sys
>>> sys.path.append('Your path to minimizer.py')
>>> import minimizer
>>> minimizer.minimize([1,2,3,4,5])
```

This computes Theorem 2.11's lower bound on the area of a five-edged centered dual two-cell with edge lengths bounded below by $(b_1, b_2, b_3, b_4, b_5)$, where $\sinh(b_i/2) = i$ for each i (see Important Note 3.5).

Finally, in subsection 3.3 we prove Proposition 3.8, on the relationship between Theorem 2.11 and Theorem 3.31 of [4], and consider some illustrative examples.

3.1. Formulas. The reduction to the trivalent case allowed by Lemma 2.10 gives a huge savings in computational expense, since there are explicit formulas for two critical quantities: the triangle area, and the circumcircle radius of a semicyclic triangle. For an arbitrary cyclic n -gon C we do not know an explicit formula in terms of side length for the area of C , and the

same is true for semicyclic circumcircle radius. Below we cite the references we know for the results we use. Any omissions are due to the author's ignorance, and additional references are welcome.

The “hyperbolic Heron formula” below was first recorded (to my knowledge) by S. Bilinski [2]. An equivalent formulation was rediscovered by W.W. Stothers [10].

Lemma 3.1. *The area of a compact hyperbolic triangle with sides of length a , b and c is:*

$$D_0(a, b, c) = 2 \cos^{-1} \left(\frac{\sinh^2(a/2) + \sinh^2(b/2) + \sinh^2(c/2) + 2}{2 \cosh(a/2) \cosh(b/2) \cosh(c/2)} \right)$$

Here D_0 refers to the area function from [5, Prop. 2.3].

Recall that we say a cyclic triangle is *semicyclic* if its longest side is also a diameter of its circumcircle or, equivalently, if its side length collection (a, b, c) lies in the space \mathcal{BC}_3 of [5, Prop. 1.11]. The “Pythagorean theorem for semicyclic hyperbolic triangles” below is recorded as Lemma 4.3 of Näätänen–Penner [9]. It can easily be derived from the hyperbolic law of sines (see eg. [1, §7.12]).

Lemma 3.2. *The circumcircle radius J of a compact, semicyclic hyperbolic triangle with shorter side lengths a and b satisfies:*

$$\sinh^2 J = \sinh^2(a/2) + \sinh^2(b/2)$$

Equivalently $b_0(a, b) = 2 \sinh^{-1} \left(\sqrt{\sinh^2(a/2) + \sinh^2(b/2)} \right)$, for b_0 as in [5, Prop. 1.12].

Corollary 3.3. *The area of a compact, semicyclic hyperbolic triangle with shorter side lengths a and b is*

$$D_0(a, b, b_0(a, b)) = 2 \sin^{-1} \left(\frac{\sinh(a/2) \sinh(b/2)}{\cosh(a/2) \cosh(b/2)} \right) = 2 \sin^{-1} (\tanh(a/2) \tanh(b/2))$$

Proof. This follows by simply substituting the formula for b_0 from Lemma 3.2 for c in the formula for D_0 from Lemma 3.1. Letting “ D_0 ” refer to $D_0(a, b, b_0(a, b))$ and “ A ” and “ B ” to $\sinh(a/2)$ and $\sinh(b/2)$, respectively, we have:

$$(3.3.1) \quad \cos(D_0/2) = \frac{2(A^2 + B^2 + 1)}{2\sqrt{A^2 + B^2 + 1} \cosh(a/2) \cosh(b/2)} = \frac{\sqrt{A^2 + B^2 + 1}}{\cosh(a/2) \cosh(b/2)}$$

Applying the identity $\sin^2 \theta = 1 - \cos^2 \theta$ and simplifying gives the result. \square

We now use the formulas above to give a self-contained, explicit statement of Theorem 2.11.

Corollary 3.4. *Let C be a compact two-cell of the centered dual complex of a locally finite set $\mathcal{S} \subset \mathbb{H}^2$ such that for some $\mathbf{b} = (b_1, \dots, b_n) \in (\mathbb{R}^+)^n$ and enumeration of the edges of C , the i^{th} edge has length at least b_i for each i . Then $\text{area}(C) \geq \min\{B_T(\mathbf{b}) \mid T \in \mathcal{T}_n, \sigma \in S_n\}$, where S_n is the symmetric group on n letters, $\sigma \in S_n$ acts on \mathbf{b} by permutation of entries,*

and \mathcal{T}_n is the collection of all compact, rooted trees T with root vertex v_T , frontier \mathcal{F} of order n , and each vertex trivalent in $T \cup \bigcup_{f \in \mathcal{F}} f$; and:

$$B_T(\mathbf{b}) = 2 \cos^{-1} \left(\frac{2 + \sum_{i=0}^2 \sinh^2(m_{e_T^i}/2)}{2 \prod_{i=0}^2 \cosh(m_{e_T^i}/2)} \right) + \sum_{v \in T^{(0)} - \{v_T\}} 2 \sin^{-1} (\tanh(b_{e_v^1}/2) \tanh(b_{e_v^2}/2))$$

Here for $v \in T^{(0)} - \{v_T\}$, e_v^1 and e_v^2 are the two edges containing v with the property that v is closer to v_T than the other endpoint. For each edge e of T , taking v to be the further endpoint of e from v_T , we recursively define $b_e = b_e(\mathbf{b})$ following [4, Lemma 3.19]:

$$b_e(\mathbf{b}) = 2 \sinh^{-1} \sqrt{\sinh^2(b_{e_v^1}/2) + \sinh^2(b_{e_v^2}/2)},$$

The three edges containing v_T are enumerated as e_T^0 , e_T^1 , and e_T^2 , and for each i , taking $i \pm 1$ modulo three, we define:

$$m_{e_T^i} = \min \left\{ b_{e_T^i}, 2 \sinh^{-1} \sqrt{\sinh^2(b_{e_T^{i+1}}/2) + \sinh^2(b_{e_T^{i-1}}/2)} \right\},$$

In particular, $m_{e_T^i} = b_{e_T^i}$ if $b_{e_T^i}$ is not maximal.

Proof. This is obtained from Theorem 2.11 by writing out the formula for $B_T(\mathbf{b})$ from Proposition 2.8 and noting that since each vertex of T has valence three in $T \cup \bigcup_{f \in \mathcal{F}} f$, Lemma 3.1 computes $D_0(m_{e_0}, b_{e_1}, \dots, b_{e_{n_T-1}})$, Lemma 3.2 computes $b_e(\mathbf{b})$ for each $e \in \mathcal{E}$, and Corollary 3.3 computes $D_0(P_v(\mathbf{b}_{\mathcal{E}}, \mathbf{b}))$ for all $v \in T^{(0)} - \{v_T\}$. \square

3.2. Programs. This section describes *minimizer.py*, a Python module containing a script *minimize()* for computing the lower bounds given by Corollary 3.4 on areas of centered dual two-cells with at most nine edges. The architecture of *minimize()* is simple, and we write it here in pseudocode:

```

define minimize(bounds)
    n = length(bounds), minlb = -1
    for tree in forest(n)
        for b in permute(bounds)
            lb = treecrawler(tree, b)
            if minlb == -1 then minlb = lb
            else minlb = min(lb, minlb)
    return minlb

```

Given an input n -tuple “bounds”, the idea is to loop over each tree in \mathcal{T}_n , and for each tree T over each permutation \mathbf{b} of bounds, computing $B_T(\mathbf{b})$ and comparing it to the minimum obtained from prior computations. Here *forest()* is a routine that produces all elements of \mathcal{T}_n for a given n ; *permute()* produces all permutations of a given tuple; and *treecrawler()* computes $B_T(\mathbf{b})$, given $T \in \mathcal{T}_n$ and an n -tuple \mathbf{b} .

Important Note 3.5. For a given n -tuple $\mathbf{b} = (b_1, \dots, b_n)$, to obtain $\min\{B_T(\sigma(\mathbf{b}))\}$ one inputs a list $[B_1, \dots, B_n]$ to *minimize()*, where $B_i = \sinh(b_i/2)$ for each i . The motivation for this choice is the nature of the explicit functions in Section 3.1.

Below we give some details on the implementations of *permute*, *treecrawler*, and *forest*.

3.2.1. *Permute*. The *itertools* Python module contains a script *permutations* that generates all permutations of a given list. Our implementation of *minimizer.py* calls this function to produce permutations of \mathbf{b} . For custom applications or use in programming languages that lack such a pre-built tool, we note that many existing permutation generation algorithms can be found, eg. in [7] or on Wikipedia. We point out one concrete example: the “Steinhaus–Johnson–Trotter algorithm”, which was later improved by S. Even, see [6].

Another thing to note is that generating all permutations of any tuple of bounds produces considerable redundancy in the output of *treecrawler*(,), on account of Remark 2.12. Guided by the KISS principle (and our limitations as a coder), we have elected not to attempt to remove this redundancy in our implementation.

3.2.2. *Treecrawler*. This function from *minimizer.py* takes two lists as input: one, “tree”, of length $n - 3$ which encodes a rooted tree T with frontier \mathcal{F} of order n , and another, “bound”, of length n which contains an entry $B_i = \sinh(b_i/2)$ for each entry b_i of a tuple \mathbf{b} of edge length bounds. A couple of observations motivate choosing this form for the inputs. First,

3.6. Every tree in \mathcal{T}_n has $n - 3$ edges.

Let \mathcal{E} be the edge set of T . Since each edge in \mathcal{E} contains two vertices of T and each edge of \mathcal{F} contains exactly one, by our trivalence hypothesis the number k of vertices of T satisfies $3k = 2|\mathcal{E}| + n$. Since T is a tree its Euler characteristic is $k - |\mathcal{E}| = 1$, so substituting gives $k = n - 2$ and $|\mathcal{E}| = n - 3$.

The second motivating observation is:

3.7. The vertices of a compact, rooted tree T with k edges and root vertex v_T can be enumerated as v_0, \dots, v_k so that for each i and j , if the arc $[v_i, v_T]$ from v_i to v_T contains v_j then $i \leq j$. Given such a numbering, enumerate the edges of T as e_0, \dots, e_{k-1} so that e_i is the initial edge of $[v_i, v_T]$ for each $i < k$. Then T is determined by the k -tuple (n_0, \dots, n_{k-1}) , where for each i , v_{n_i} is the nearer vertex of e_i to v_T .

One may produce the desired enumeration of the vertices of T by first listing all those at maximal distance d from v_T in T , then listing those at distance $d - 1$, and so forth. Note that any such enumeration has $v_T = v_k$. And since T is a tree there is a *unique* arc joining v_i to v_T for all $i < k$, so e_i is uniquely determined for each such i by the requirement above. This yields k unique edges; all of them, since T is a tree with $k + 1$ vertices.

Figure 3.1 depicts the rooted trees with one to three edges, with vertices enumerated and the resulting encoding tuples following 3.7.

Treecrawler distributes the bounds of “bound” to frontier edges of T in the opposite order from that of the vertices. That is, having enumerated the vertices of T as v_0, \dots, v_{n-3} following 3.7, we enumerate \mathcal{F} as $\{f_1, \dots, f_n\}$ so that if $i < j$ then $n_i \geq n_j$, where v_{n_i} and




FIGURE 3.1. The rooted trees with $2 \leq k \leq 4$ vertices, with root vertices enlarged and vertices labeled, and the associated $(k - 1)$ -tuples.

v_{n_j} are the vertices of T contained f_i and f_j , respectively. Then we bound d_{f_i} below by b_i for each i . (This choice is made because it is computationally least expensive.)

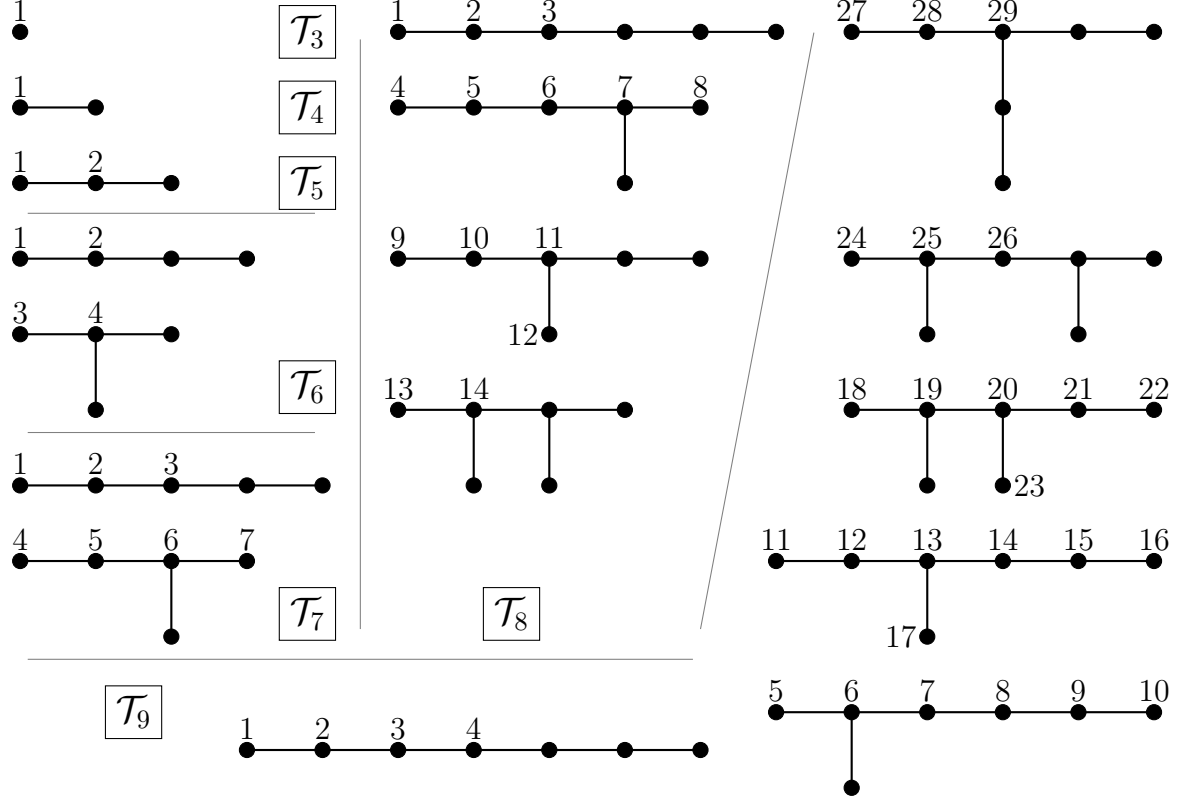
The idea of the program is to recursively compute b_{e_i} and $D_0(P_{v_i}(\mathbf{b}))$ using the formulas of Lemmas 3.1 and 3.2. Having done so for a given i , we append b_{e_i} to a list “edgelengths” containing the lengths e_j , and we add $D_0(P_{v_i}(\mathbf{b}))$ to the number “totalarea” that records the sum of the $D_0(P_{v_j}(\mathbf{b}))$, for $j < i$. The main observation here is that for each $i < n - 3$, v_i is contained in exactly two edges of $\mathcal{E} \cup \mathcal{F}$ not equal to e_i , and if either of these is of the form $e_j \in \mathcal{E}$ then $j < i$.

3.2.3. Forest. This is a library of all elements of \mathcal{T}_n for $3 \leq n \leq 9$. Recall that an element T of \mathcal{T}_n is a compact, rooted tree with frontier \mathcal{F} of order n , such that each vertex of T has valence three in $T \cup \bigcup_{f \in \mathcal{F}} f$, and therefore valence at most three in T . In *forest.txt* and Figure 3.2, which depicts its members, we track only the internal structure of T — the idea being that each vertex of T gets as many frontier edges as necessary to bring its valence up to three in $T \cup \bigcup_{f \in \mathcal{F}} f$.

By 3.6 above, each $T \in \mathcal{T}_n$ has $k = n - 3$ edges, so *forest.txt* encodes T by a string of length k following the scheme in 3.7. The first line of *forest.txt* is a key: a (0-based) list L whose k th entry $L[k]$ records the number of lines down that the codes for trees with k edges begin. In fact the codes begin one line below that: the line $L[k]$ below the first contains a single integer, which is the total number of codes for k -edged trees. Each tree code occupies a single line. It is the code described in 3.7 but written in reverse order (*minimizer.py* un-reverses the order when reading the code).

Figure 3.2 depicts the $T \in \mathcal{T}_n$ for $3 \leq n \leq 9$. For each n , trees are numbered in the figure according to their position in the corresponding list in *forest.txt*: if a vertex v of a tree T with $n - 3$ edges there is numbered i then the element of \mathcal{T}_n represented by T with root vertex v is also represented by the i th code from the top of the list in *forest.txt* containing codes for \mathcal{T}_n . For instance, for the vertex labeled 14 on the five-edged tree T in the figure, enumerating the vertices of T following 3.7 yields the five-tuple $(2, 2, 5, 5, 5)$. The reverse of this is the fourteenth and final 5-edged tree code in *forest.txt*.

A vertex is not numbered if and only if it is equivalent to one that is under a non-trivial automorphism of the tree that contains it. Note that the vertex numbering prescribed in 3.7


 FIGURE 3.2. The collections \mathcal{T}_n for $n \leq 9$.

is not necessarily unique, so each code is one of possibly several representing the same tree. In the case above the other possibilities are $(3, 3, 5, 5, 5)$ and $(4, 4, 5, 5, 5)$.

We record a couple of observations to support the classification of trees in Figure 3.2. A tree T with k edges has diameter $d \leq k$. Let γ be an edge arc of length equal to d (such an arc is the horizontal part of each tree in the Figure). The $k - d$ edges not contained in γ lie in a disjoint union of maximal subtrees of T that do not contain edges of γ . If T_0 is such a subtree then its intersection point v_0 with γ can be at most $d - d_0$ away from each boundary vertex of γ , where d_0 is the maximal distance in T_0 from v_0 to another vertex. Otherwise the diameter hypothesis would be violated.

It is now just a matter of enumerating the possible $d \leq k$, and for each d , the possibilities for T_0 . We note also that d cannot be too small. For instance, in the six-edge case if d were equal to 3 then three edges of T would lie outside γ . There are only two possible attachment points for the trees T_0 , the interior vertices of γ . So one such T_0 must contain at least two edges, hence it must have $d_0 \geq 2$. But this contradicts our assumption on d , since each interior vertex of γ is at distance 2 from one of its boundary vertices.

3.3. Examples. Here we analyze a couple of examples that illustrate important basic features of the bounds of Theorem 2.11. Before the first, we pause to observe:

Proposition 3.8. *For any $n > 4$ and $d > 0$ the bound $\min\{B_T(\sigma(\mathbf{b}))\}$ given by Theorem 2.11 for $\mathbf{b} = (d, \dots, d) \in \mathbb{R}^n$ is larger than the bound $(n - 2)A_m(d)$ from [4, Theorem 3.31]. For $n = 3$ and $n = 4$ the two results offer equal bounds.*

Proof. We first touch on the case $n = 3$. In this special case the lower bound of [4, Thrm. 3.31] is the area of an equilateral triangle with all side lengths d ; ie. $D_0(d, d, d)$ in the language of [5, L. 2.1]. The only tree in \mathcal{T}_3 is $T = \{v_T\}$, and chasing the definition of B_T in Proposition 2.8 gives $B_T(d, d, d) = D_0(d, d, d)$. We therefore assume below that $n \geq 4$.

For the given \mathbf{b} we note that $\sigma(\mathbf{b}) = \mathbf{b}$ for all $\sigma \in S_n$, so in the bound of Theorem 2.11 the minimum need only be taken over $T \in \mathcal{T}_n$. By 3.6 above, each $T \in \mathcal{T}_n$ has $n - 2$ vertices. Thus applying the definition of B_T , the result will follow by observing for any such T that $D_0(P_v(\mathbf{b}_\mathcal{E}, \mathbf{b})) \geq A_m(d)$ for each $v \in T^{(0)} - \{v_T\}$; that $D_0(m_{e_0}, b_{e_1}, b_{e_2}) \geq A_m(d)$, where e_0, e_1 and e_2 in $\mathcal{E} \cup \mathcal{F}$ contain v_T with b_{e_0} maximal among the b_{e_i} ; and that strict inequality holds for some vertex if $n > 4$.

The tuple $\mathbf{b}_\mathcal{E} \in (\mathbb{R}^+)^{\mathcal{E}}$ referenced above is defined in Proposition 2.8 by $b_e = b_e(\mathbf{b})$ for each $e \in \mathcal{E}$, where $b_e: (\mathbb{R}^+)^{\mathcal{F}} \rightarrow \mathbb{R}^+$ is from [4, Lemma 3.19]. For each $v \in T^{(0)} - \{v_T\}$, taking e_v as the initial edge of the arc joining v to v_T , $b_{e_v} = b_0(b_{e_1}, b_{e_2})$ by the definition of b_e in [4, L. 3.19], where b_0 is the function from [5, Prop. 1.12] and $e_1, e_2 \in \mathcal{E} \cup \mathcal{F}$ are the other two edges containing v . Thus $D_0(P_v(\mathbf{b}_\mathcal{E}, \mathbf{b})) = D_0(b_{e_1}, b_{e_2}, b_0(b_{e_1}, b_{e_2}))$.

For v and e_1, e_2 as above, if $e_i \in \mathcal{F}$ then $b_{e_i} = d$ by our hypothesis here. It follows from Property (1) of [4, L. 3.19] that $b_{e_i} > d$ if $e_i \in \mathcal{E}$. Corollary 2.4 and Proposition 1.12 of [5] together imply that if $x \leq x'$ and $y \leq y'$ then $D_0(x, y, b_0(x, y)) \leq D_0(x', y', b_0(x', y'))$, and that strict inequality holds here if $x < x'$ or $y < y'$. Therefore $D_0(P_v(\mathbf{b}_\mathcal{E}, \mathbf{b})) \geq D_0(d, d, b_0(d, d))$ for all $v \in T^{(0)} - \{v_T\}$, and the inequality is strict unless v has valence one in T .

Property (1) of [4, L. 3.19] can be used here to show the stronger fact that $b_e \geq b_0(d, d)$ for each $e \in \mathcal{E}$. By its definition in Proposition 2.8 and the monotonicity of b_0 ([5, Prop. 1.12]) we also have $m_{e_0} \geq b_0(d, d)$ for e_0, e_1 and e_2 containing v_T as above. Therefore again $D_0(m_{e_0}, b_{e_1}, b_{e_2}) \geq D_0(d, d, b_0(d, d))$ by [5, Cor. 2.3], with strict inequality if v_T does not have valence one in T .

As pointed out in the statement of [4, Thrm. 3.31], $A_m(d)$ is the area of a semicyclic triangle with two sides of length d . But $b_0: (\mathbb{R}^+)^2 \rightarrow \mathbb{R}^+$ is characterized by the property that $(x, y, b_0(x, y)) \in \mathcal{BC}_3$ has unique largest entry $b_0(x, y)$ for all x and y , where \mathcal{BC}_3 is the space parametrizing semicyclic triangles, so $A_m(d) = D_0(d, d, b_0(d, d))$. It therefore follows from above that $\min\{B_T(\sigma(\mathbf{b}))\} \geq (n - 2)A_m(d)$, and that strict inequality holds unless each vertex of T has valence one in T . This in turn only holds if T has two vertices and one edge; ie. if $n = 4$.

It is now straightforward to verify in the case $n = 4$ for the lone rooted tree T with one edge, and $\mathbf{b} = (d, d, d, d)$, that $B_T(\mathbf{b}) = 2D_0(d, d, b_0(d, d)) = 2A_m(d)$. \square

Example 3.9. Here we consider the n -tuple \mathbf{b}_n with all entries $b = 2\sinh^{-1}(1)$ (ie. with $\sinh(b/2) = 1$) and allow n to vary between 4 and 9. The outputs of Theorem 3.31 of [4]




FIGURE 3.3. The two cases of Example 3.10.

and Theorem 2.11 here are recorded in the table below, in each case truncated after three decimal places. Recall from Important Note 3.5 that the input for *minimizer.minimize()* corresponding to \mathbf{b} is $[1, \dots, 1]$.

The “tree number” line below refers to the number in Figure 3.2 (or *forest.txt*) of the tree realizing $\min\{B_T(\mathbf{b})\}$.

n	4	5	6	7	8	9
$(n-2)A_m(b)$	2.094	3.141	4.188	5.235	6.283	7.330
$\min\{B_T(\mathbf{b}_n)\}$	2.094	3.295	4.526	5.818	7.107	8.441
tree number	1	2	4	6	14	20

 TABLE 1. The outputs of [4, Thrm. 3.31] and Theorem 2.11 on \mathbf{b}_n as above.

It is not hard to show using Corollary 3.3 that for our chosen b , $A_m(b) = \pi/3$. So the exact value of $(n-2)A_m(b)$ is $2\pi/3$ for $n = 4$, π for $n = 5$, and so on. Note that the difference between bounds grows monotonically with n .

A takeaway from the proof of Proposition 3.8 is that for a particular T , with $\mathbf{b} = (1, \dots, 1)$, vertices with valence one in T contribute the least to $B_T(\mathbf{b})$. It is therefore not surprising that for each n the rooted tree realizing $\min\{B_T(\mathbf{b}_n)\}$ has the maximum possible number of vertices of valence one among all trees in \mathcal{T}_n . Note moreover that the root vertex of each such example has maximal valence.

Example 3.10. Now we explicitly analyze the simplest case with unequal entries, taking $\mathbf{c} = (b, b, b, x)$ with $\sinh(b/2) = 1$ as in the previous example. There is only one tree $T \in \mathcal{T}_4$, and following the scheme of *treecrawler* we label its frontier as $\{f_0, f_1, f_2, f_3\}$ so that f_0 and f_1 terminate at the root vertex v_T and f_2 and f_3 at the other one. Since \mathbf{c} has many identical entries, by Remark 2.12 there are only two permutations of \mathbf{c} to check: the identity, which assigns x to f_3 and b to all others, and a permutation σ that assigns x to f_0 . These are pictured in Figure 3.3.

Frontier edges are dashed in the figure, and each such edge f_i is labeled with $\sinh(c_{f_i}/2)$, where c_f is the corresponding entry of \mathbf{c} (or $\sigma(\mathbf{c})$ in the other case). In particular, $X = \sinh(x/2)$. The edge e of T points toward v_T (which hence is the right vertex in each case). It is labeled with $b_e = b_e(\mathbf{c})$ as prescribed in Corollary 3.4.

In order to compute $B_T(\mathbf{c})$ and $B_T(\sigma(\mathbf{c}))$ we must first determine the m_{e_i} as in Corollary 3.4, where e_0, e_1 and e_2 are the edges containing v_T . We will take $e_0 = f_0$, $e_1 = f_1$, and $e_2 = e$

in each case. Then in the identity case, $\sinh(m_{e_2}/2) = \sqrt{X^2 + 1}$ if $X \leq 1$ and otherwise $\sinh(m_{e_2}/2) = \sqrt{2}$, and $m_{e_i} = b$ for $i = 0, 1$ and any x . This gives:

$$B_T(\mathbf{c}) = \begin{cases} 2 \sin^{-1} \left(\frac{X}{\sqrt{2(X^2+1)}} \right) + 2 \cos^{-1} \left(\frac{5+X^2}{4\sqrt{X^2+2}} \right) & X \leq 1 \\ 2 \sin^{-1} \left(\frac{X}{\sqrt{2(X^2+1)}} \right) + 2 \sin^{-1} \left(\frac{1}{2} \right) & X > 1 \end{cases}$$

In each case above the contribution from v_T is the second summand. For $X > 1$ we simplify by using the observation from the proof of Lemma 3.3, that for inputs in \mathcal{BC}_3 , Heron's formula reduces to the one from that result.

In case σ we have $m_{e_0} = x$, $m_{e_1} = b$ and $\sinh(m_{e_2}/2) = \sqrt{X^2 + 1}$ if $X \leq 1$; $m_{e_0} = x$, $m_{e_1} = b$ and $\sinh(m_{e_2}/2) = \sqrt{2}$ for $1 < X \leq \sqrt{3}$; and $\sinh(m_{e_0}/2) = \sqrt{3}$, $m_{e_1} = b$ and $\sinh(m_{e_2}/2) = \sqrt{2}$ if $X > \sqrt{3}$. Applying this gives:

$$B_T(\sigma(\mathbf{c})) = \begin{cases} 2 \sin^{-1} \left(\frac{1}{2} \right) + 2 \sin^{-1} \left(\frac{X}{\sqrt{2(X^2+1)}} \right) & X \leq 1 \\ 2 \sin^{-1} \left(\frac{1}{2} \right) + 2 \cos^{-1} \left(\frac{5+X^2}{2\sqrt{6(X^2+1)}} \right) & 1 < X \leq \sqrt{3} \\ 2 \sin^{-1} \left(\frac{1}{2} \right) + 2 \sin^{-1} \left(\frac{1}{\sqrt{3}} \right) & X > \sqrt{3} \end{cases}$$

Again the contribution from v_T is the second summand. To understand which of $B_T(\mathbf{c})$ or $B_T(\sigma(\mathbf{c}))$ is larger for various values of x , it helps to look at them in a different way. Below to save space we replace a by $\sinh(a/2)$ for any entry a of an input to the area function D_0 .

$$\begin{aligned} B_T(\sigma(\mathbf{c})) - B_T(\mathbf{c}) &= D_0(1, 1, \sqrt{2}) - D_0(1, X, \sqrt{2}) > 0 & \text{for } X \leq 1 \\ B_T(\mathbf{c}) - B_T(\sigma(\mathbf{c})) &= D_0(1, X, \sqrt{X^2 + 1}) - D_0(1, X, \sqrt{2}) > 0 & \text{for } 1 < X \leq \sqrt{3} \\ B_T(\mathbf{d}) - B_T(\sigma(\mathbf{c})) &= D_0(1, X, \sqrt{X^2 + 1}) - D_0(1, \sqrt{3}, \sqrt{2}) > 0 & \text{for } X > \sqrt{3} \end{aligned}$$

The inequalities in each case follow from Proposition 2.3 of [5]. Therefore as long as $X \leq 1$, $\min\{B_T(\mathbf{c}), B_T(\sigma(\mathbf{c}))\} = B_T(\mathbf{c})$, and otherwise it is $B_T(\sigma(\mathbf{c}))$.

Note that the given bound is constant for $x \geq 2 \sinh^{-1}(\sqrt{3})$. Its value, truncated after three decimal places, is 2.278 for such x . This may be compared with the bound of 2.094 given for $\mathbf{b}_4 = (b, b, b, b)$ in Example 1.

Remark 3.11. The cases of Example 3.10 where $X \leq \sqrt{3}$ are distinguished by the fact that taking $\mathbf{b}_{\mathcal{F}} = \mathbf{c}$ or $\mathbf{b}_{\mathcal{F}} = \sigma(\mathbf{c})$ as appropriate to minimize B_T , the tuple $\mathbf{b}_{\mathcal{E}}$ produced by Proposition 2.8 has $(\mathbf{b}_{\mathcal{E}}, \mathbf{b}_{\mathcal{F}}) \in \overline{Ad}_T(\mathbf{b})$. Another way of saying this, again using the notation of Proposition 2.8, is that $m_{e_0} = b_{e_0}$.

We say that an arbitrary n -tuple \mathbf{b} with the analogous property is “geometrically reasonable”. That is, for $T \in \mathcal{T}_n$ and $\sigma \in S_n$ such that $B_T(\sigma(\mathbf{b}))$ realizes the bounds of Theorem 2.11, we should have $(\mathbf{b}_{\mathcal{E}}, \sigma(\mathbf{b})) \in \overline{Ad}_T(\sigma(\mathbf{b}))$. The reason is that in this case there is a metric triangle complex \mathfrak{T} with combinatorics prescribed by T and geometry by $(\mathbf{b}_{\mathcal{E}}, \sigma(\mathbf{b}))$, and a map $\mathfrak{T} \rightarrow \mathbb{H}^2$ whose image we expect in many cases to be a centered dual cell with edge length collection $\sigma(\mathbf{b})$ and area equalling the bound.

The idea here is that \mathfrak{T} has a triangular face for each $v \in T^{(0)}$ which is a hyperbolic triangle with edge length collection $P_v(\mathbf{b}_\mathcal{E}, \sigma(\mathbf{b}))$, and for each edge e of T the faces of \mathfrak{T} corresponding to its endpoints are glued along their sides with length b_e . A map $\mathfrak{T} \rightarrow \mathbb{H}^2$ is determined by choosing an isometric embedding of $P_{v_T}(\mathbf{b}_\mathcal{E}, \sigma(\mathbf{b}))$ and analytically continuing outward from v_T in T , forcing the restriction to each $P_v(\mathbf{b}_\mathcal{E}, \sigma(\mathbf{b}))$ to be an isometry. This map is therefore a local isometry on the interior of \mathfrak{T} .

Such a map may fail to be an isometry if there is branching at the vertices, or if different arms of T determine regions of \mathfrak{T} with overlapping images. Even in the absence of these, a vertex of some P_v may end up inside the circumcircle of some other P_w . But for many geometrically reasonable \mathbf{b} these pathologies will not occur, and for these the image of \mathfrak{T} will be a centered dual cell of the set of images \mathcal{S} of its vertices.

REFERENCES

- [1] Alan F. Beardon. *The geometry of discrete groups*, volume 91 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1983.
- [2] Stanko Bilinski. Zur Begründung der elementaren Inhaltslehre in der hyperbolischen Ebene. *Math. Ann.*, 180:256–268, 1969.
- [3] Jason DeBlois. The Delaunay tessellation in hyperbolic space. Preprint. arXiv:1308.4899.
- [4] Jason DeBlois. The centered dual and the maximal injectivity radius of hyperbolic surfaces. *Geom. Topol.*, 19(2):953–1014, 2015.
- [5] Jason DeBlois. The geometry of cyclic hyperbolic polygons. *Rocky Mountain J. Math.*, 46(3):801–862, 2016.
- [6] Wikipedia entry: The Steinhaus–Johnson–Trotter algorithm. en.wikipedia.org/wiki/SteinhausJohnsonTrotter_algorithm.
- [7] Donald E. Knuth. *The art of computer programming. Vol. 1*. Addison-Wesley, Reading, MA, 1997. Fundamental algorithms, Third edition.
- [8] Sadayoshi Kojima and Yosuke Miyamoto. The smallest hyperbolic 3-manifolds with totally geodesic boundary. *J. Differential Geom.*, 34(1):175–192, 1991.
- [9] M. Nääätänen and R. C. Penner. The convex hull construction for compact surfaces and the Dirichlet polygon. *Bull. London Math. Soc.*, 23(6):568–574, 1991.
- [10] Wilson Stothers’ Geometry Pages. www.maths.gla.ac.uk/wws/cabripages/cabri0.html. In particular, see <http://www.maths.gla.ac.uk/wws/cabripages/hyperbolic/overview.html>.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF PITTSBURGH, 301 THACKERAY HALL, PITTSBURGH, PA 15260

E-mail address: jdeblois@pitt.edu